MIT Technology Review

The politics issue

Vol 121 No. 5
Sep/Oct 2018
$9.99 USD
$10.99 CAD

Technology is threatening our democracy. How do we save it?

The "neuropolitics" consultants who hack voters' brains

Two very different experiments in digital government

Science fiction: Karl Schroeder on the future of fake news


Artificial intelligence is changing every business.
Don’t be left behind.

MIT Technology Review

EmTech Digital

March 25–26, 2019
St. Regis Hotel
San Francisco, CA

EmTechDigital.com/2019

Gideon Lichfield is editor in chief of MIT Technology Review.

"Big Data Will Save Politics." When we put those words on the cover of *MIT Technology Review* in 2013, Barack Obama had just won reelection with the help of a crack team of data scientists and engineers. The Arab Spring had already cooled into a grim Arab Winter, but the social-media platforms that had powered the uprisings were still basking in the afterglow. Silicon Valley was full of hope and hubris about its power to democratize the world.

Today, with Cambridge Analytica, fake news, election hacking, and the shrill cacophony that dominates social media, technology feels as likely to destroy politics as to save it. The tech firms and their boosters either didn't imagine that "democratizing" technologies would be used by anti-democrats too, or else believed that truth and freedom would inevitably defeat misinformation and repression.

Our cover in January 2013

Those aren't the only reasons to worry that tech is threatening democracy. One can wonder if voters are even capable of making free choices anymore, as political campaigns microtarget citizens in key districts with increasingly surgical precision (see Alex Howard's story on page 28) and "neuropolitical" consultants use cognitive science to identify preferences people didn't even know they had (Elizabeth Svoboda, page 64). Convincing video "deepfakes" are becoming easy to create, as Will Knight found by making one himself (page 36). Twitter bots that currently spew cookie-cutter propaganda on autopilot will soon be smart enough to engage individual users in simulated conversation, as bot expert Lisa-Maria Neudert explains (page 72). Big data is disrupting the cozy world of political lobbying, as Andrew Zaleski shows in his profile of FiscalNote founder Tim Hwang (page 56)—but, not surprisingly, it may be deep-pocketed companies that benefit the most. And with November's midterms approaching, the US's election systems are still woefully insecure, as Martin Giles reports (page 44).

In this issue we also take a close look at the famed "filter bubble" effect that's blamed for political polarization. The striking data visualizations from John Kelly and Camille François on page 22 show that the effect is real: like-minded people band together online, and extremism shouts loudest. Yet contrary to popular belief, you encounter more opposing views on the internet, not fewer, Adam Piore reports (page 18). What happens, as Zeynep Tufekci explains in the opening essay (page 10), is that online communities create a tribal mind-set that makes people less receptive to these views.

Finally, as Tufekci also argues, it would be as much of a mistake to simply blame technology for democracy's ills as it was to hail technology as democracy's savior. The turmoil of the past few years is of our societies' own making. We need to pay more attention to how technology exacerbates those problems, but fixing them is not solely, or even mainly, a technological challenge.

Yet while tech isn't the solution, it can be part of it. We end this issue on two hopeful notes. Chris Horton reports from Taiwan on an experiment in digital democracy that shows how simple design rules for an online platform can lead people toward consensus instead of division. And in Karl Schroeder's science fiction story ... but I won't give away the plot. I hope this issue gives you plenty to chew on. Let me know what you think at gideon.lichfield@technologyreview.com.

NEPHI NIVEN

TECHONOMY

Techonomy 2018 | November 11-13 | The Ritz-Carlton Half Moon Bay

Where leaders explore the intersection of tech, business and society


Arati Prabhakar


Bernard Tyson


Padmasree Warrior


Marc Benioff

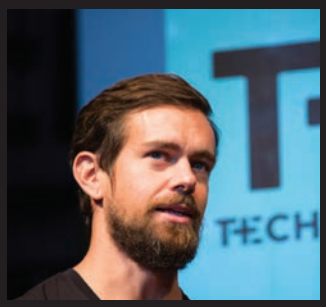
Jack Dorsey


Stacy Brown-Philpot


Mark Zuckerberg


James Park


Patrick Collison


John Chambers


Jessica Rosenworcel


Sen. Cory Booker


Marissa Mayer


Peter Thiel


Jaron Lanier

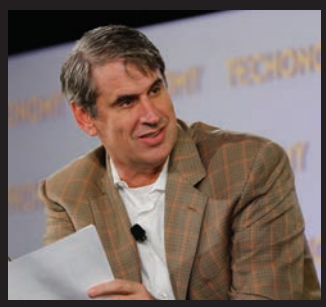
Bill Gurley

Join a community from business, academia, government and the media for a carefully curated, interactive agenda at the annual Techonomy retreat. Our program is imbued with optimism, though we'll also tackle this moment of skepticism, as tech retrenches. We discuss it all at the beautiful, cliffside Ritz-Carlton in Half Moon Bay, California.

For more information and to request your invitation, go to **techonomy.com/techreview**

Techonomy 2018 Partners Include

accenture ERICSSON Johnson & Johnson PHILIPS standard W2O

The politics issue

How we got here

On the cover:
Illustration by Harry Campbell

Where we are now

What comes next

Note: In our last issue we used the term "techonomy" on this contents page, and on the cover and in several locations in the magazine, to represent the intersection of technology and the economy. We did not intend to imply any affiliation with or endorsement by New York–based Techonomy Media Inc., which produces content and events that underscore tech's impact on business and society. For more information on Techonomy Media and its November conference in California, visit www.techonomy.com.

No two cancers are alike. The same goes for cancer treatments. Innovative immunotherapies like CAR-T can now reprogram patients' immune systems to destroy the disease. Fighting cancer has never been more personal. **This is the future of medicine. For all of us.**

GoBoldly.com

GOBOLDLY

America's Biopharmaceutical Companies®

Customer service and subscription inquiries

National
800-877-5230

International
903-636-1115

E-mail
customer_service@
mittechnologyreview.info

Web
www.technologyreview.com/
customerservice

MIT Records (alums only)
617-253-8270

Reprints
techreview@wrightsmedia.com
877-652-5295

Licensing and permissions
licensing@technologyreview.com

MIT Technology Review

One Main Street, 13th Floor
Cambridge, MA 02142
Tel: 617-475-8000

The mission of *MIT Technology Review* is to bring about better-informed and more conscious decisions about technology through authoritative, influential, and trustworthy journalism.

Get MIT Technology Review delivered to your inbox.

THE DOWNLOAD
Your daily dose of what's up in emerging technology

CLOCKING IN
A daily look at the workplace of the future

CHAIN LETTER
Blockchains, cryptocurrencies, and why they matter

THE ALGORITHM
News and views on the latest in artificial intelligence

Stay in the know:
technologyreview.com/inbox

MIT Technology Review

Foundational inventions that change entire industries.

At Qualcomm, inventing comes first. When we connected the phone to the internet, our foundational inventions created the mobile revolution. Now, as we lead the world to 5G, our inventions are going to enable new industries to be created, and the next great product the world can't live without.

qualcomm.com/weinvent

Qualcomm

Inventing the tech the world loves

How we got here

How the internet went from being the darling of the Arab Spring to everybody's punching bag (p10). "Filter bubbles" don't work the way you think they do (p18). Twitter data shows what political divides really look like (p22). How US election campaign tech has evolved and where it's going (p28). And a taxonomy of former internet boosters and their excuses (p32).

1

The road from Tahrir to Trump

To understand how digital technologies went from instruments for spreading democracy to weapons for attacking it, you have to look beyond the technologies themselves.

By ZEYNEP TUFEKCI

1. The euphoria of discovery

As the Arab Spring convulsed the Middle East in 2011 and authoritarian leaders toppled one after another, I traveled the region to try to understand the role that technology was playing. I chatted with protesters in cafés near Tahrir Square in Cairo, and many asserted that as long as they had the internet and the smartphone, they would prevail. In Tunisia, emboldened activists showed me how they had used open-source tools to track the shopping trips to Paris that their autocratic president's wife had taken on government planes. Even Syrians I met in Beirut were still optimistic; their country had not yet descended into a hellish war. The young people had energy, smarts, humor, and smartphones, and we expected that the region's fate would turn in favor of their democratic demands.

Egyptian protesters asserted that as long as they had the internet and the smartphone, they would prevail.

Back in the United States, at a conference talk in 2012, I used a screenshot from a viral video recorded during the Iranian street protests of 2009 to illustrate how the new technologies were making it harder for traditional information gatekeepers—like governments and the media—to stifle or control dissident speech. It was a difficult image to see: a young woman lay bleeding to death on the sidewalk. But therein resided its power. Just a decade earlier, it would most likely never have been taken (who carried video cameras all the time?), let alone gone viral (how, unless you owned a TV station or a newspaper?). Even if a news photographer had happened to be there, most news organizations wouldn't have shown such a graphic image.

At that conference, I talked about the role of social media in breaking down what social scientists call "pluralistic ignorance"—the belief that one is alone in one's views when in reality everyone has been collectively silenced. That, I said, was why social media had fomented so much rebellion: people who were previously isolated in their dissent found and drew strength from one another.

Twitter, the company, retweeted my talk in a call for job applicants to "join the flock." The implicit understanding was that Twitter was a force for good in the world, on the side of the people and their revolutions. The new information gatekeepers, which didn't see themselves as gatekeepers but merely as neutral "platforms," nonetheless liked the upending potential of their technologies.

I shared in the optimism. I myself hailed from the Middle East and had been watching dissidents use digital tools to challenge government after government.

But a shift was already in the air.

During the Tahrir uprising, Egypt's weary autocrat, Hosni Mubarak, had clumsily cut off internet and cellular service. The move backfired: it restricted the flow of information coming out of Tahrir Square but caused international attention on Egypt to spike. He hadn't understood that in the 21st century it is the flow of attention, not information (which we already have too much of), that matters. Besides, friends of the spunky Cairo revolutionaries promptly flew in with satellite phones, allowing them to continue giving interviews and sending images to global news organizations that now had even more interest in them.

Within a few weeks, Mubarak was forced out. A military council replaced him. What it did then foreshadowed much of what was to come. Egypt's Supreme Council of the Armed Forces promptly opened a Facebook page and made it the exclusive outlet for its communiqués. It had learned from Mubarak's mistakes; it would play ball on the dissidents' turf.

Within a few years, Egypt's online sphere would change dramatically. "We had more influence when it was just us on Twitter," one activist prominent on social media told me. "Now it is full of bickering between dissidents [who are] being harassed by government supporters." In 2013, on the heels of protests against a fledgling but divisive civilian government, the military would seize control.

Power always learns, and powerful tools always fall into its hands. This is a hard lesson of history but a solid one. It is key to understanding how, in seven years, digital technologies have gone from being hailed as tools of freedom and change to being blamed for upheavals in Western democracies—for enabling increased polarization, rising authoritarianism, and meddling in national elections by Russia and others.

But to fully understand what has happened, we also need to examine how human social dynamics, ubiquitous digital connectivity, and the business models of tech giants combine to create an environment where misinformation thrives and even true information can confuse and paralyze rather than informing and illuminating.

2. The audacity of hope

Barack Obama's election in 2008 as the first African-American president of the United States had prefigured the Arab Spring's narrative of technology empowering the underdog. He was an unlikely candidate who had emerged triumphant, beating first Hillary Clinton in the Democratic primary and then

ED GILES/GETTY IMAGES

his Republican opponent in the general election. Both his 2008 and 2012 victories prompted floods of laudatory articles on his campaign's tech-savvy, data-heavy use of social media, voter profiling, and microtargeting. After his second win, *MIT Technology Review* featured Bono on its cover, with the headline "Big Data Will Save Politics" and a quote: "The mobile phone, the Net, and the spread of information—a deadly combination for dictators."

However, I and many others who watched authoritarian regimes were already worried. A key issue for me was how microtargeting, especially on Facebook, could be used to wreak havoc with the public sphere. It was true that social media let dissidents know they were not alone, but online microtargeting could also create a world in which you wouldn't know what messages your neighbors were getting or how the ones aimed at you were being tailored to your desires and vulnerabilities.

Digital platforms allowed communities to gather and form in new ways, but they also dispersed existing communities, those that had watched the same TV news and read the same newspapers. Even living on the same street meant less when information was disseminated through algorithms designed to maximize revenue by keeping people glued to screens. It was a shift from a public, collective politics to a more private, scattered one, with political actors collecting more and more personal data to figure out how to push just the right buttons, person by person and out of sight.

All this, I feared, could be a recipe for misinformation and polarization.

Shortly after the 2012 election, I wrote an op-ed for the *New York Times* voicing these worries. Not wanting to sound like a curmudgeon, I understated my fears. I merely advocated transparency and accountability for political ads and content on social media, similar to systems in place for regulated mediums like TV and radio.

The backlash was swift. Ethan Roeder, the data director for the Obama 2012 campaign, wrote a piece headlined "I Am Not Big Brother," calling such worries "malarkey." Almost all the data scientists and Democrats I talked to were terribly irritated by my idea that technology could be anything but positive. Readers who commented on my op-ed thought I was just being a spoilsport. Here was a technology that allowed Democrats to be better at elections. How could this be a problem?

3. The illusion of immunity

The Tahrir revolutionaries and the supporters of the US Democratic Party weren't alone in thinking they would always have the upper hand.

The US National Security Agency had an arsenal of hacking tools based on vulnerabilities in digital technologies—bugs, secret backdoors, exploits, shortcuts in the (very advanced) math, and massive computing power. These tools were dubbed "nobody but us" (or NOBUS, in the acronym-loving intelligence community), meaning no one else could exploit them, so there was no need to patch the vulnerabilities or

There were laudatory articles about Barack Obama's use of voter profiling and microtargeting.

The generals in Egypt learned from Hosni Mubarak's mistakes.

ALEX WONG/GETTY IMAGES; PETER MACDIARMID/GETTY IMAGES

The NSA had an arsenal of hacking tools dubbed NOBUS.

make computer security stronger in general. The NSA seemed to believe that weak security online hurt its adversaries a lot more than it hurt the NSA.

That confidence didn't seem unjustified to many. After all, the internet is mostly an American creation; its biggest companies were founded in the United States. Computer scientists from around the world still flock to the country, hoping to work for Silicon Valley. And the NSA has a giant budget and, reportedly, thousands of the world's best hackers and mathematicians.

Since it's all classified, we cannot know the full story, but between 2012 and 2016 there was at least no readily visible effort to significantly "harden" the digital infrastructure of the US. Nor were loud alarms raised about what a technology that crossed borders might mean. Global information flows facilitated by global platforms meant that someone could now sit in an office in Macedonia or in the suburbs of Moscow or St. Petersburg and, for instance, build what appeared to be a local news outlet in Detroit or Pittsburgh.

There doesn't seem to have been a major realization within the US's institutions—its intelligence agencies, its bureaucracy, its electoral machinery—that true digital security required both better technical infrastructure and better public awareness about the risks of hacking, meddling, misinformation, and more. The US's corporate dominance and its technical wizardry in some areas seemed to have blinded the country to the brewing weaknesses in other, more consequential ones.

4. The power of the platforms

In that context, the handful of giant US social-media platforms seem to have been left to deal as they saw fit with what problems might emerge. Unsurprisingly, they prioritized their stock prices and profitability. Throughout the years of the Obama administration, these platforms grew boisterously and were essentially unregulated. They spent their time solidifying their technical chops for deeply surveilling their users, so as to make advertising on the platforms ever more efficacious. In less than a decade, Google and Facebook became a virtual duopoly in the digital ad market.

Facebook also gobbled up would-be competitors like WhatsApp and Instagram without tripping antitrust alarms. All this gave it more data, helping it improve its algorithms for keeping users on the platform and targeting them with ads. Upload a list of already identified targets and Facebook's AI engine will helpfully find much bigger "look-alike" audiences that may be receptive to a given message. After 2016, the grave harm this feature could do would become obvious.

Meanwhile, Google—whose search rankings can make or break a company, service, or politician, and whose e-mail service had a billion users by 2016—also operated the video platform YouTube, increasingly a channel for information and propaganda around the world. A *Wall Street Journal* investigation earlier this year found that YouTube's recommendation algorithm tended to drive viewers toward extremist content by suggesting edgier versions of whatever they were watching—a good way to hold their attention.

This was lucrative for YouTube but also a boon for conspiracy theorists, since people are drawn to novel and shocking claims. "Three degrees of Alex Jones" became a running joke: no matter where you started on YouTube, it was said, you were never more than three recommendations away from a video by the right-wing conspiracist who popularized the idea that the Sandy Hook school shooting in 2012 had never happened and the bereaved parents were mere actors playing parts in a murky conspiracy against gun owners.

Though smaller than Facebook and Google, Twitter played an outsize role thanks to its popularity among journalists and politically engaged people. Its open philosophy and easygoing approach to pseudonyms suits rebels around the world, but it also appeals to anonymous trolls who hurl abuse at women, dissidents, and minorities. Only earlier this year did it crack down on the use of bot accounts that trolls used to automate and amplify abusive tweeting.

Twitter's pithy, rapid-fire format also suits anyone with a professional or instinctual understanding of attention, the crucial resource of the digital economy.

Say, someone like a reality TV star. Someone with an uncanny ability to come up with belittling, viral nicknames for his opponents, and to make boastful promises that resonated with a realignment in

American politics—a realignment mostly missed by both Republican and Democratic power brokers.

Donald Trump, as is widely acknowledged, excels at using Twitter to capture attention. But his campaign also excelled at using Facebook as it was designed to be used by advertisers, testing messages on hundreds of thousands of people and microtargeting them with the ones that worked best. Facebook had embedded its own employees within the Trump campaign to help it use the platform effectively (and thus spend a lot of money on it), but they were also impressed by how well Trump himself performed. In later internal memos, reportedly, Facebook would dub the Trump campaign an "innovator" that it might learn from. Facebook also offered its services to Hillary Clinton's campaign, but it chose to use them much less than Trump's did.

Digital tools have figured significantly in political upheavals around the world in the past few years, including others that left elites stunned: Britain's vote to leave the European Union, and the far right's gains in Germany, Hungary, Sweden, Poland, France, and elsewhere. Facebook helped Philippine strongman Rodrigo Duterte with his election strategy and was even cited in a UN report as having contributed to the ethnic-cleansing campaign against the Rohingya minority in Myanmar.

However, social media isn't the only seemingly democratizing technology that extremists and authoritarians have co-opted. Russian operatives looking to hack into the communications of Democratic Party officials used Bitcoin—a cryptocurrency founded to give people anonymity and freedom from reliance on financial institutions—to buy tools such as virtual private networks, which can help one cover one's traces online. They then used these tools to set up fake local news organizations on social media across the US.

There they started posting materials aimed at fomenting polarization. The Russian trolls posed as American Muslims with terrorist sympathies and as white supremacists who opposed immigration. They posed as Black Lives Matter activists exposing police brutality and as people who wanted to acquire guns to shoot police officers. In so doing, they not only fanned the flames of division but provided those in each group with evidence that their imagined opponents were indeed as horrible as they suspected. These trolls also incessantly harassed journalists and Clinton supporters online, resulting in a flurry of news stories about the topic and fueling a (self-fulfilling) narrative of polarization among the Democrats.

Donald Trump's campaign excelled at using Facebook as it was designed to be used by advertisers.

BRETT CARLSEN/STRINGER/GETTY IMAGES

5. The lessons of the era

How did all this happen? How did digital technologies go from empowering citizens and toppling dictators to being used as tools of oppression and discord? There are several key lessons.

First, the weakening of old-style information gatekeepers (such as media, NGOs, and government and academic institutions), while empowering the underdogs, has also, in another way, deeply disempowered underdogs. Dissidents can more easily circumvent censorship, but the public sphere they can now reach is often too noisy and confusing for them to have an impact. Those hoping to make positive social change have to convince people both that something in the world needs changing and there is a constructive, reasonable way to change it. Authoritarians and extremists, on the other hand, often merely have to muddy the waters and weaken trust in general so that everyone is too fractured and paralyzed to act. The old gatekeepers blocked some truth and dissent, but they blocked many forms of misinformation too.

Second, the new, algorithmic gatekeepers aren't merely (as they like to believe) neutral conduits for both truth and falsehood. They make their money by keeping people on their sites and apps; that aligns their incentives closely with those who stoke outrage, spread misinformation, and appeal to people's existing biases and preferences. Old gatekeepers failed in many ways, and no doubt that failure helped fuel mistrust and doubt; but the new gatekeepers *succeed*

by fueling mistrust and doubt, as long as the clicks keep coming.

Third, the loss of gatekeepers has been especially severe in local journalism. While some big US media outlets have managed (so far) to survive the upheaval wrought by the internet, this upending has almost completely broken local newspapers, and it has hurt the industry in many other countries. That has opened fertile ground for misinformation. It has also meant less investigation of and accountability for those who exercise power, especially at the local level. The Russian operatives who created fake local media brands across the US either understood the hunger for local news or just lucked into this strategy. Without local checks and balances, local corruption grows and trickles up to feed a global corruption wave playing a major part in many of the current political crises.

The fourth lesson has to do with the much-touted issue of filter bubbles or echo chambers—the claim that online, we encounter only views similar to our own. This isn't completely true. While algorithms will often feed people some of what they already want to hear, research shows that we probably encounter a wider variety of opinions online than we do offline, or than we did before the advent of digital tools.

Rather, the problem is that when we encounter opposing views in the age and context of social media, it's not like reading them in a newspaper while sitting alone. It's like hearing them from the opposing team while sitting with our fellow fans in a football stadium. Online, we're connected with our communities, and we seek approval from our like-minded peers. We bond with our team by yelling at the fans of the other one. In sociology terms, we strengthen our feeling of "in-group" belonging by increasing our distance from and tension with the "out-group"—us versus them. Our cognitive universe isn't an echo chamber, but our social one is. This is why the various projects for fact-checking claims in the news, while valuable, don't convince people. Belonging is stronger than facts.

A similar dynamic played a role in the aftermath of the Arab Spring. The revolutionaries were caught up in infighting on social media as they broke into ever smaller groups, while at the same time authoritarians were mobilizing their own supporters to attack the dissidents, defining them as traitors or foreigners. Such "patriotic" trolling and harassment is probably more common, and a bigger threat to dissidents, than attacks orchestrated by governments.

This is also how Russian operatives fueled polarization in the United States, posing simultaneously as immigrants and white supremacists, angry Trump

The old information gatekeepers blocked some truth and dissent but also many forms of misinformation.

CHIP SOMODEVILLA/GETTY IMAGES

supporters and "Bernie bros." The content of the argument didn't matter; they were looking to paralyze and polarize rather than convince. Without old-style gatekeepers in the way, their messages could reach anyone, and with digital analytics at their fingertips, they could hone those messages just like any advertiser or political campaign.

Fifth, and finally, Russia exploited the US's weak digital security—its "nobody but us" mind-set—to subvert the public debate around the 2016 election. The hacking and release of e-mails from the Democratic National Committee and the account of Clinton campaign manager John Podesta amounted to a censorship campaign, flooding conventional media channels with mostly irrelevant content. As the Clinton e-mail scandal dominated the news cycle, neither Trump's nor Clinton's campaign got the kind of media scrutiny it deserved.

This shows, ultimately, that "nobody but us" depended on a mistaken interpretation of what digital security means. The US may well still have the deepest offensive capabilities in cybersecurity. But Podesta fell for a phishing e-mail, the simplest form of hacking, and the US media fell for attention hacking. Through their hunger for clicks and eyeballs, and their failure to understand how the new digital sphere operates, they were diverted from their core job into a confusing swamp. Security isn't just about who has more Cray supercomputers and cryptography experts but about understanding how attention, information overload, and social bonding work in the digital era.

This potent combination explains why, since the Arab Spring, authoritarianism and misinformation have thrived, and a free-flowing contest of ideas has not. Perhaps the simplest statement of the problem, though, is encapsulated in Facebook's original mission statement (which the social network changed in 2017, after a backlash against its role in spreading misinformation). It was to make the world "more open and connected." It turns out that this isn't necessarily an unalloyed good. Open to *what*, and connected *how*? The need to ask those questions is perhaps the biggest lesson of all.

6. The way forward

What is to be done? There are no easy answers. More important, there are no purely digital answers.

There are certainly steps to be taken in the digital realm. The weak antitrust environment that allowed a few giant companies to become near-monopolies should be reversed. However, merely breaking up these giants without changing the rules of the game online may simply produce a lot of smaller companies that use the same predatory techniques of data surveillance, microtargeting, and "nudging."

There are no easy answers, and no purely digital answers.

Ubiquitous digital surveillance should simply end in its current form. There is no justifiable reason to allow so many companies to accumulate so much data on so many people. Inviting users to "click here to agree" to vague, hard-to-pin-down terms of use doesn't produce "informed consent." If, two or three decades ago, before we sleepwalked into this world, a corporation had suggested so much reckless data collection as a business model, we would have been horrified.

There are many ways to operate digital services without siphoning up so much personal data. Advertisers have lived without it before, they can do so again, and it's probably better if politicians can't do it so easily. Ads can be attached to content, rather than directed to people: it's fine to advertise scuba gear to me if I am on a divers' discussion board, for example, rather than using my behavior on other sites to figure out that I'm a diver and then following me around everywhere I go—online or offline.

But we didn't get where we are simply because of digital technologies. The Russian government may have used online platforms to remotely meddle in US elections, but Russia did not create the conditions of social distrust, weak institutions, and detached elites that made the US vulnerable to that kind of meddling.

Russia meddled in US politics, but it didn't create the conditions that made the US vulnerable to such meddling.

CHRIS MCGRATH/GETTY IMAGES

Digital connectivity provided the spark, but the kindling was everywhere.

PETER MACDIARMID/GETTY IMAGES

Russia did not make the US (and its allies) initiate and then terribly mishandle a major war in the Middle East, the after-effects of which—among them the current refugee crisis—are still wreaking havoc, and for which practically nobody has been held responsible. Russia did not create the 2008 financial collapse: that happened through corrupt practices that greatly enriched financial institutions, after which all the culpable parties walked away unscathed, often even richer, while millions of Americans lost their jobs and were unable to replace them with equally good ones.

Russia did not instigate the moves that have reduced Americans' trust in health authorities, environmental agencies, and other regulators. Russia did not create the revolving door between Congress and the lobbying firms that employ ex-politicians at handsome salaries. Russia did not defund higher education in the United States. Russia did not create the global network of tax havens in which big corporations and the rich can pile up enormous wealth while basic government services get cut.

These are the fault lines along which a few memes can play an outsize role. And not just Russian memes: whatever Russia may have done, domestic actors in the United States and Western Europe have been eager, and much bigger, participants in using digital platforms to spread viral misinformation.

Even the free-for-all environment in which these digital platforms have operated for so long can be seen as a symptom of the broader problem, a world in which the powerful have few restraints on their actions while everyone else gets squeezed. Real wages in the US and Europe are stuck and have been for decades while corporate profits have stayed high and taxes on the rich have fallen. Young people juggle multiple, often mediocre jobs, yet find it increasingly hard to take the traditional wealth-building step of buying their own home—unless they already come from privilege and inherit large sums.

If digital connectivity provided the spark, it ignited because the kindling was already everywhere. The way forward is not to cultivate nostalgia for the old-world information gatekeepers or for the idealism of the Arab Spring. It's to figure out how our institutions, our checks and balances, and our societal safeguards should function in the 21st century—not just for digital technologies but for politics and the economy in general. This responsibility isn't on Russia, or solely on Facebook or Google or Twitter. It's on us.

Zeynep Tufekci is an associate professor at the University of North Carolina and a contributing opinion writer at the New York Times.

No, big tech didn't make us

Last fall, Deb Roy, one of the US's foremost experts on social media, attended a series of roundtables in small towns in middle America—places like Platteville, Wisconsin, and Anamosa, Iowa. It wasn't what Roy, who runs the Laboratory for Social Machines at the MIT Media Lab, was used to: there were no computer screens in the rooms, no tweets or posts to examine. Instead, he just listened to community leaders and local residents talk, face to face, about their neighbors.

What he heard alarmed him greatly.

"'I found out what they said on Facebook,'" Roy recalls one elderly woman saying. "'Their views are so extreme and so unacceptable to me that I no longer see the point in engaging with them.'" It was a sentiment he heard again and again.

"These are people they see on a regular basis in their small towns," he says. "They used to agree to disagree. When divisiveness and balkanization are reflected at the hyperlocal level—where even when we have access to one another, the digital realm is actually silencing our speech and cutting us off from one another in the physical realm—something is profoundly wrong."

In 2014, Roy set up his MIT lab specifically to study, among other things, how social media could be used to help break through the partisan arguing that typically divides people. He may be uniquely positioned to make such an attempt. From 2013 to 2017, the Canadian-born engineer served as "chief media scientist" for Twitter, collecting and analyzing social-media chatter. When he opened his lab, Twitter not only granted him full access to "the firehose"—every single tweet ever produced, in real time—but ponied up $10 million to help him make sense of all this information about people's interests, preferences, and activities, and find ways to use it for public benefit.

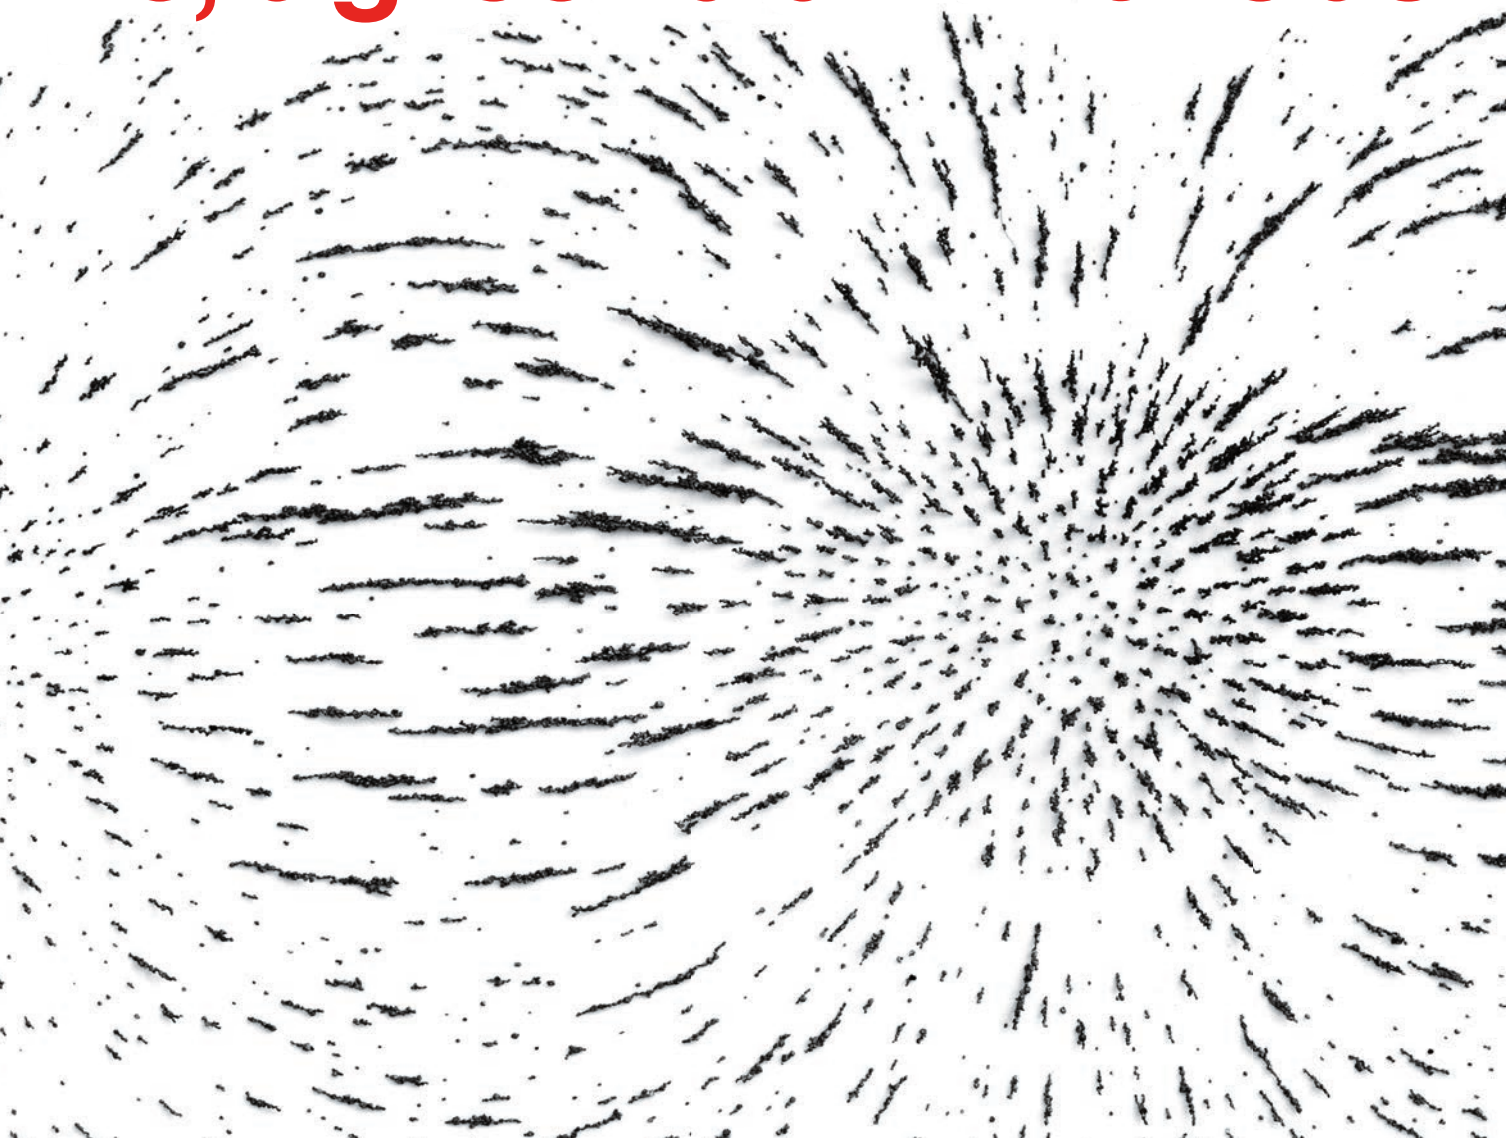

polarized (but it sure helps)

Social-media bubbles tell only part of the story of why we're so divided. The rest is in our heads.

By Adam Piore

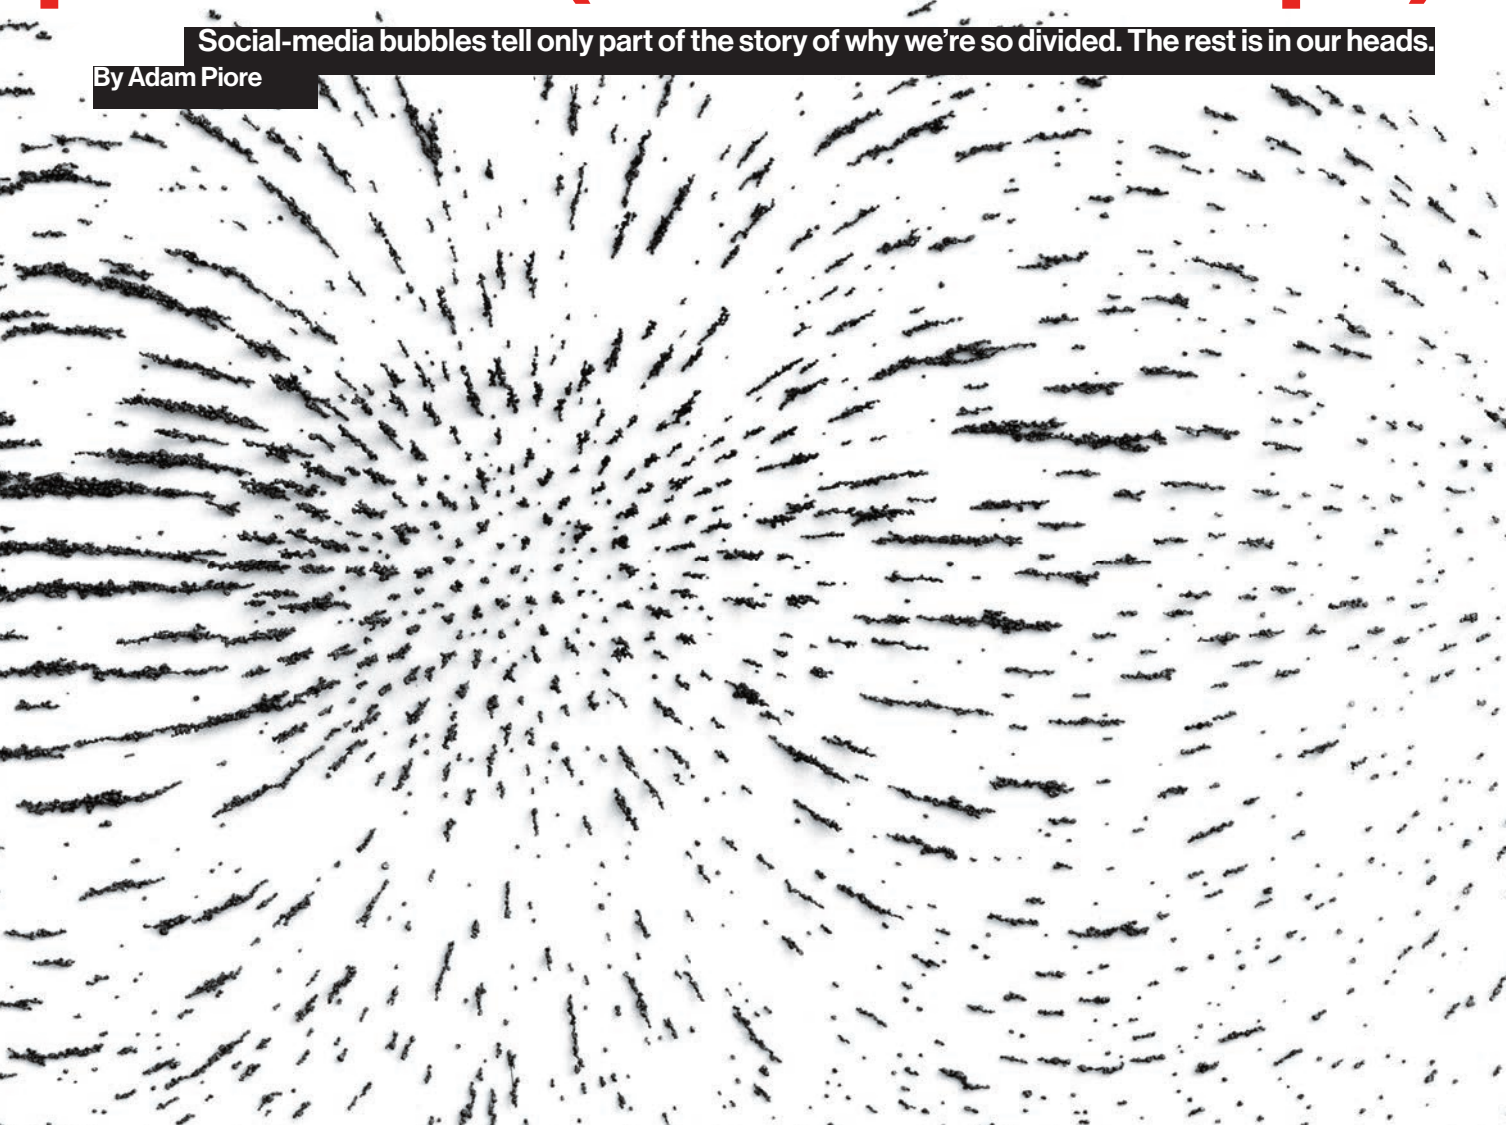

For Roy and a number of other researchers who study the internet's impact on society, the most concerning problem highlighted by the 2016 election isn't that the Russians used Twitter and Facebook to spread propaganda, or that the political consulting firm Cambridge Analytica illicitly gained access to the private information of more than 50 million Facebook users. It's that we have all, quite voluntarily, retreated into hyperpartisan virtual corners, owing in no small part to social media and internet companies that determine what we see by monitoring what we have clicked on in the past and giving us more of the same. In the process, opposing perspectives are sifted out, and we're left with content that reinforces what we already believe.

This is the famous "filter bubble," a concept popularized in the 2011 book of the same name by Eli Pariser, an internet activist and founder of the viral video site Upworthy. "Ultimately, democracy works only if we citizens are capable of thinking beyond our narrow self-interest," wrote Pariser. "But to do so, we need a shared view of the world we coinhabit. The filter bubble pushes us in the opposite direction—it creates the impression that our narrow self-interest is all that exists."

Or does it? The research suggests that things are not quite that simple.

Some kind of war

The legal scholar Cass Sunstein warned way back in 2007 that the internet was giving rise to an "era of enclaves and niches." He cited a 2005 experiment in Colorado in which 60 Americans from conservative Colorado Springs and liberal Boulder, two

cities about 100 miles apart, were assembled into small groups and asked to deliberate on three controversial issues (affirmative action, gay marriage, and an international treaty on global warming). In almost every case, people held more extreme positions after they spoke with like-minded others.

"The Internet makes it exceedingly easy for people to replicate the Colorado experiment online, whether or not that is what they are trying to do," Sunstein wrote in the *Chronicle of Higher Education*. "There is a general risk that those who flock together, on the Internet or elsewhere, will end up both confident and wrong, simply because they have not been sufficiently exposed to counterarguments. They may even think of their fellow citizens as opponents or adversaries in some kind of 'war.'"

But is social media really at fault here? In a study published earlier this year in *Proceedings of the National Academy of Sciences*, researchers at Stanford University examined political polarization in the US and found that it was increasing far faster among the demographic groups *least* likely to use social media and the internet. "The 65-year-olds are polarizing more quickly than the younger age group, which is the opposite of what you'd expect if social media and the internet were the driver," says Levi Boxell, the lead author of the study.

What's more, most people aren't as stuck in echo chambers as some would have us think, according to Grant Blank, a research fellow at the Oxford Internet Institute, and collaborators who surveyed adults in the UK and Canada.

"We have five different ways in which the echo chamber could be defined, and it really doesn't matter which one you use because the results are very consistent across all of them—there is no echo chamber," Blank says. "People actually read lots of media. They consume, on average, five different media sources, about three offline and two online, and they encounter diverse opinions. People encounter things they disagree with, and they change their mind based on things that they encounter in media."

Even Pariser, who gave the filter bubble its name, agrees the internet isn't entirely to blame. It might explain why liberal elites didn't see Trump coming, since a large portion of middle America was absent from liberals' social-media feeds: indeed, Blank's work concluded that most researchers finding such an effect were studying only these cultural elites. But for most Trump supporters, talk radio, local news, and Fox—a pre-internet filter bubble—were far more important sources than tweets or fake news on Facebook.

Data from the polling firm Pew backs up the idea that polarization doesn't come just from the internet. After the 2016 election, Pew found that 62 percent of Americans got news from social-media sites, but—in a parenthetical ignored in most articles about the study—only 18 percent said they did so "often." A more recent Pew study found that only about 5 percent said they had "a lot" of trust in the information.

"The internet is absolutely not the causal factor here," says Ethan Zuckerman, who directs MIT's Center for Civic Media. "But I think we're experiencing a phenomenon that began with Fox News and now is sort of extending into the social-media space."

Right. So what, if anything, can we do about it?

Three attempts at a fix

After the 2016 election, Zuckerman and some collaborators designed a tool called Gobo that lets people adjust their own bubbles via sliders that control content filters. For instance, the "politics" slider ranges from "my perspective" to "lots of perspectives." Choosing the latter end exposes people to media outlets they probably wouldn't normally see.

Facebook, however, showed little interest in adopting Gobo. "What Facebook is worried about is that they believe—and they're probably right—that very few people would actually want to diversify their feed," Zuckerman says.

Another tool, Social Mirror, was developed by members of Deb Roy's lab. Earlier this year they reported on the results of an experiment conducted with the tool, which uses data visualization to give Twitter users a bird's-eye view of how their network of followers and friends fits into the overall universe of Twitter. Most of those recruited to use the tool were politically active Twitter users, and many were surprised to learn just how cocooned inside far-right or far-left bubbles they were.

The impact of the experiment was short-lived, however. Though a week after it ended some participants were following a more diverse set of Twitter accounts than before, two to three weeks later most had gone back to homogeneity. And in another twist, people who ended up following more contrarian accounts—suggested by the researchers to help them diversify their Twitter feeds—subsequently reported that they'd be even *less* inclined to talk to people with opposing political views.

Lousy results such as this have led Zuckerman toward a more radical idea for countering filter bubbles: the creation of a taxpayer-funded social-media platform with a civic mission to provide a "diverse and global view of the world."

The early United States, he noted in an essay for the *Atlantic*, featured a highly partisan press tailored to very specific audiences. But publishers and editors for the most part abided by a strong cultural norm, republishing a wide range of stories from different parts of the nation and reflecting different political leanings. Public broadcasters in many democracies have also focused on providing a wide range of perspectives. It's not realistic, Zuckerman argues, to expect the same from outlets like Facebook: their business model drives them to pander to our natural human desire to congregate with others like ourselves.

A public social-media platform with a civic mission, says Zuckerman, could push unfamiliar perspectives into our feeds and push us out of our comfort zones. Scholars could review

PREVIOUS SPREAD: ISTOCK

algorithms to make sure we're seeing an unbiased representation of views. And yes, he admits, people would complain about publicly funding such a platform and question its even-handedness. But given the lack of other viable solutions, he says, it's worth a shot.

The problem is us

Jay Van Bavel, a social psychologist at New York University, has studied social posts and analyzed which ones are most likely to gain traction. He's found that "group identification" posts activate the most primitive non-intellectual parts of the brain. So, for example, if a Republican politician tells people that immigrants are moving in and changing the culture or taking locals' jobs, or if a Democrat tells female students that Christian activists want to ban women's rights, their words have power. Bavel's research suggests that if you want to overcome partisan divisions, avoid the intellect and focus on the emotions.

After the Social Mirror experiment, members of Roy's lab debuted a project called FlipFeed, which exposed people on Twitter to others with different political views. Martin Saveski, the study's lead author, says the point was to change how people felt about the other side. One of the experiments prompted participants to imagine, whenever they came across an opposing view, that they were disagreeing with a friend. Those given this prompt were more likely to say they would like to speak with the person in the future, and that they understood why the other person held an opposing view.

The results were congruent with another observation made by Pariser. He's noticed that some of the best political discussions online happen in sports forums, where people are already united by the common love of a team. The assumption there is that all are fans of the team first, and conservative or liberal second. There's an emotional connection before politics even enters the discussion.

If you look at all the various projects from Zuckerman and Roy and others, what they're really trying to do is employ technology to get us to engage with content outside our political bubbles. But is that workable? As Roy himself says, "I don't think there are any pure, simply technological fixes."

Maybe in the end it's up to us to decide to expose ourselves to content from a wider array of sources, and then to engage with it. Sound unappealing? Well, consider the alternative: your latest outraged political post didn't accomplish much, because the research shows that anyone who read it almost certainly agreed with you already.

Adam Piore is the author of The Body Builders: Inside the Science of the Engineered Human, published in 2017.

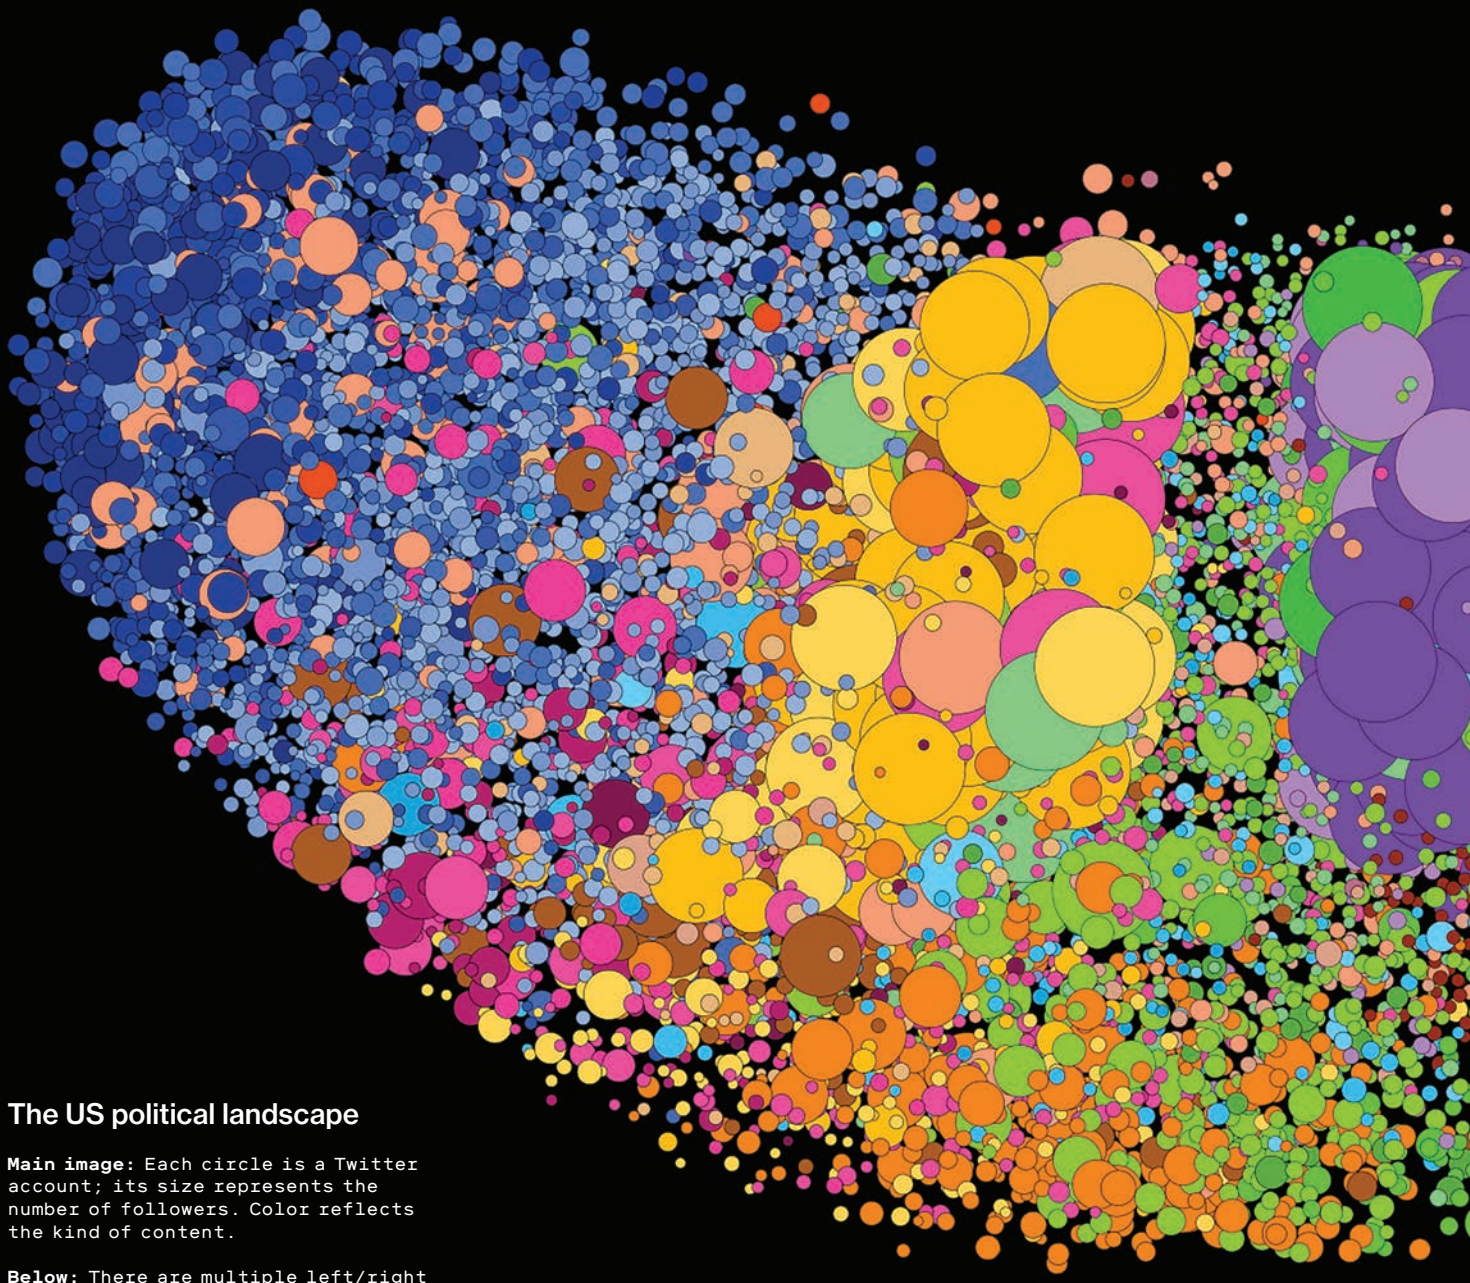

The US political landscape

Main image: Each circle is a Twitter account; its size represents the number of followers. Color reflects the kind of content.

Below: There are multiple left/right groups, with the political parties closer to the center and anti-/pro-Trump groups at the extremes.

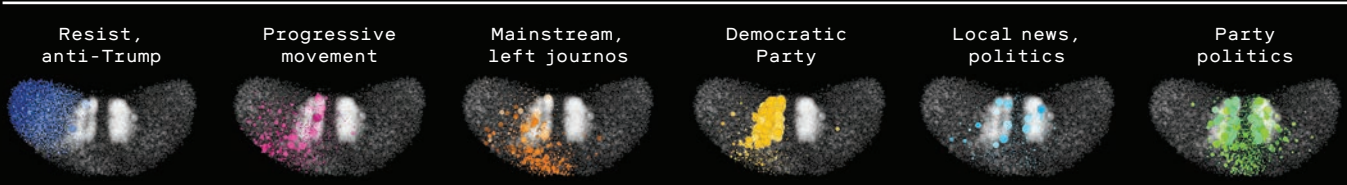

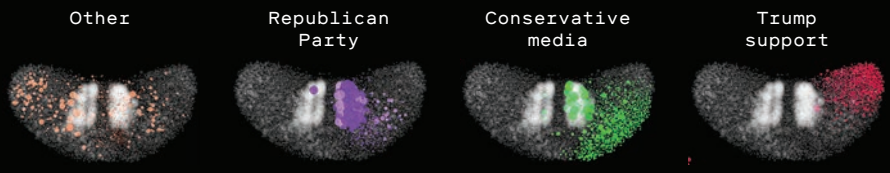

A vision of division

Maps of Twitter activity show how political polarization manifests online and why divides are so hard to bridge.

By John Kelly and Camille François

American public life has become increasingly ideologically segregated as newspapers have given way to screens. But societies have experienced extremism and fragmentation without the assistance of Silicon Valley for centuries. And the polarization in the US began long ago, with the rise of 24-hour cable news. So just how responsible is the internet for today's divisions? And are they really as bad as they seem?

In this Twitter map of the US political landscape, accounts that follow one another are clustered together, and they are color-coded by the kinds of content they commonly share. At first glance, it might seem reassuring: although there are clear echo chambers, there is also an intertwined network of elected officials, the press, and political and policy professionals. There are extremes, but they are mediated through a robust middle.

However, as the diagrams on the following pages will show, that middle is a lot weaker than it looks, and this makes public discourse vulnerable both to extremists at home and to manipulation by outside actors such as Russia.

Noisy, partisan bots

The center of the political universe is far quieter than the polarized wings. This plot of average daily tweets (vertical axis) from the network shown on the previous page shows that the extreme partisans on both sides are screaming while the center whispers. It also shows divisions being amplified by bots on both sides: we see clearly automated activity, with accounts churning out a hundred tweets a day or more on a common schedule. Hundreds of accounts (especially on the left) have identical daily tweet counts, further evidence that they are bots.

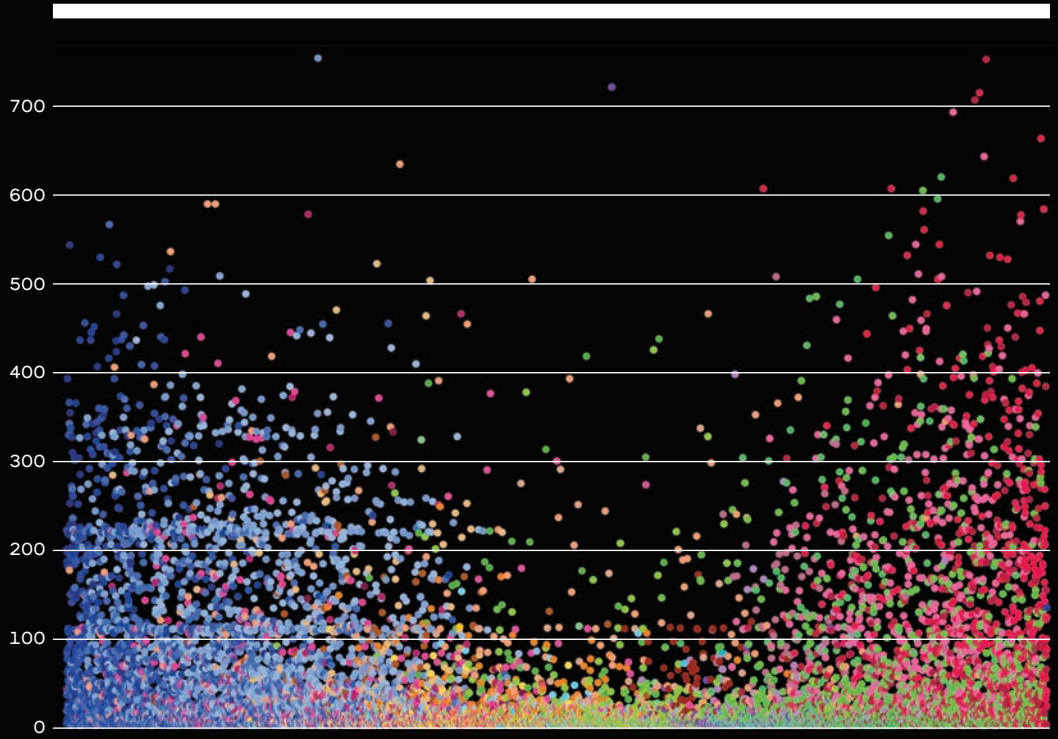

The silence of the center

The polarization looks even more extreme when the accounts are plotted according to their "valence," a measure of how politically homogeneous their connections are. A valence of 0 means an account follows or is followed only by progressive accounts, while 1 means it's connected only to conservative accounts.

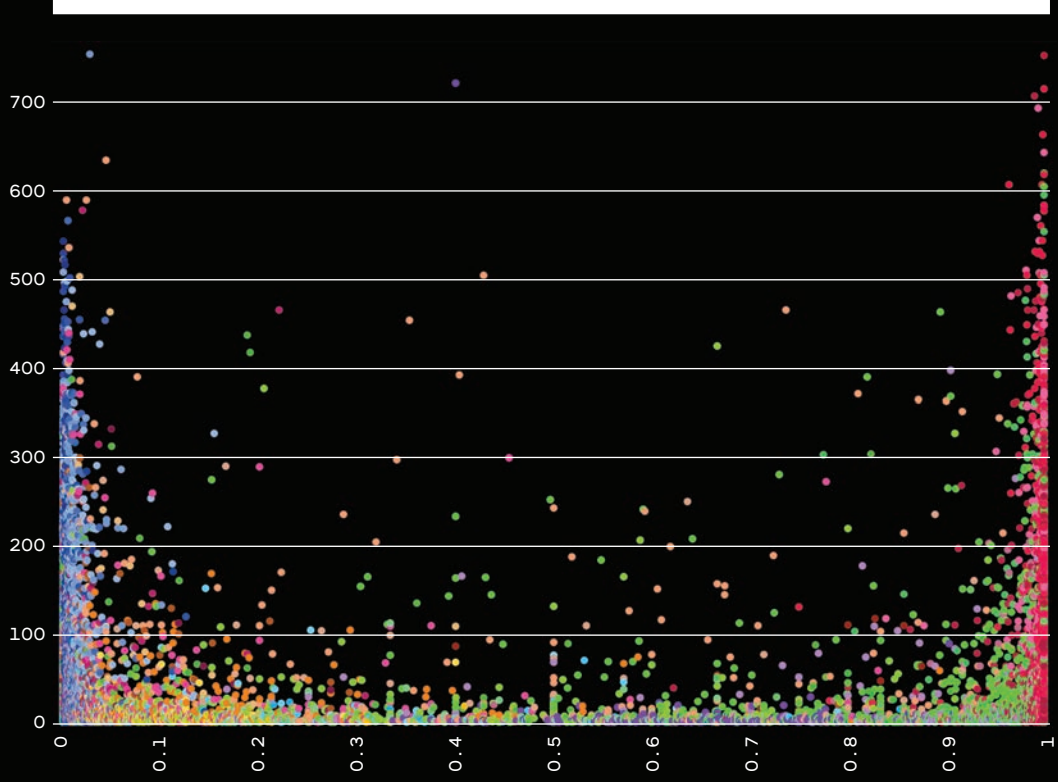

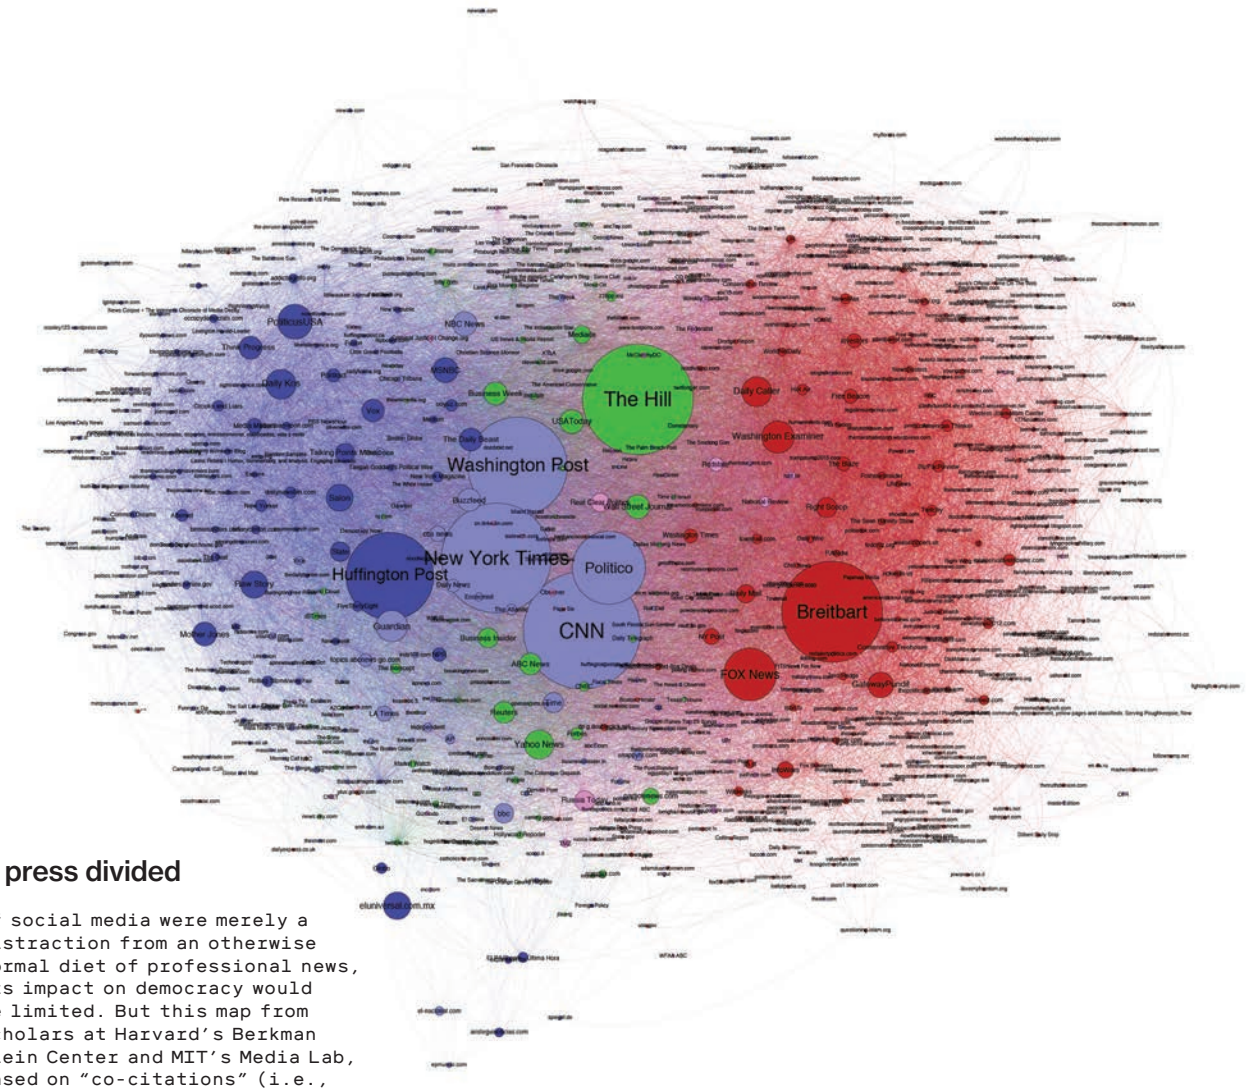

A press divided

If social media were merely a distraction from an otherwise normal diet of professional news, its impact on democracy would be limited. But this map from scholars at Harvard's Berkman Klein Center and MIT's Media Lab, based on "co-citations" (i.e., who links to whom), shows that the media world is bifurcated too.

This plot shows which news sources are cited the most by the Twitter accounts in the network map on pages 22-23. Traditional mainstream journalistic sources are cited primarily by the left, while the right is served by sources such as Fox, Breitbart, and True Pundit.

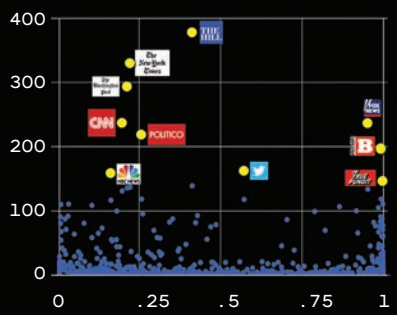

If we analyze the same data by looking at citations of individual articles rather than news sources, the divide is even more stark. The articles that get the most tweets represent the most partisan views on both left and right.

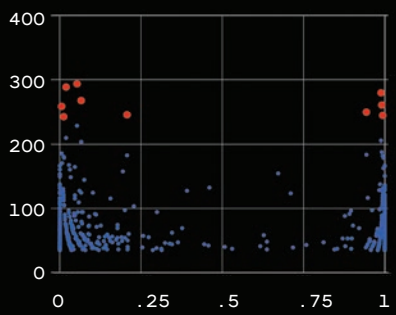

How Russian trolls exploit division

The polarization we saw on the previous pages is fertile ground for misinformation operations such as the one Russia conducted to influence the 2016 US election. Instead of trying to force their messages into the mainstream, these adversaries target polarized communities and "embed" fake accounts within them. The false personas engage with real people in those communities to build credibility. Once their influence has been established, they can introduce new viewpoints and amplify divisive and inflammatory narratives that are already circulating. It's the digital equivalent of moving to an isolated and tight-knit community, using its own language quirks and catering to its obsessions, running for mayor, and then using that position to influence national politics.

The first of these two maps shows the US political spectrum on the eve of the 2016 election. The second highlights the followers of a 30-something American woman called Jenna Abrams, a following gained with her viral tweets about slavery, segregation, Donald Trump, and Kim Kardashian. Her far-right views endeared her to conservatives, and her entertaining shock tactics won her attention from several mainstream media outlets and got her into public spats with prominent people on Twitter, including a former US ambassador to Russia. Her following in the right-wing Twittersphere enabled her to influence the broader political conversation. In reality, she was one of many fake personas created by the infamous St. Petersburg troll farm known as the Internet Research Agency.

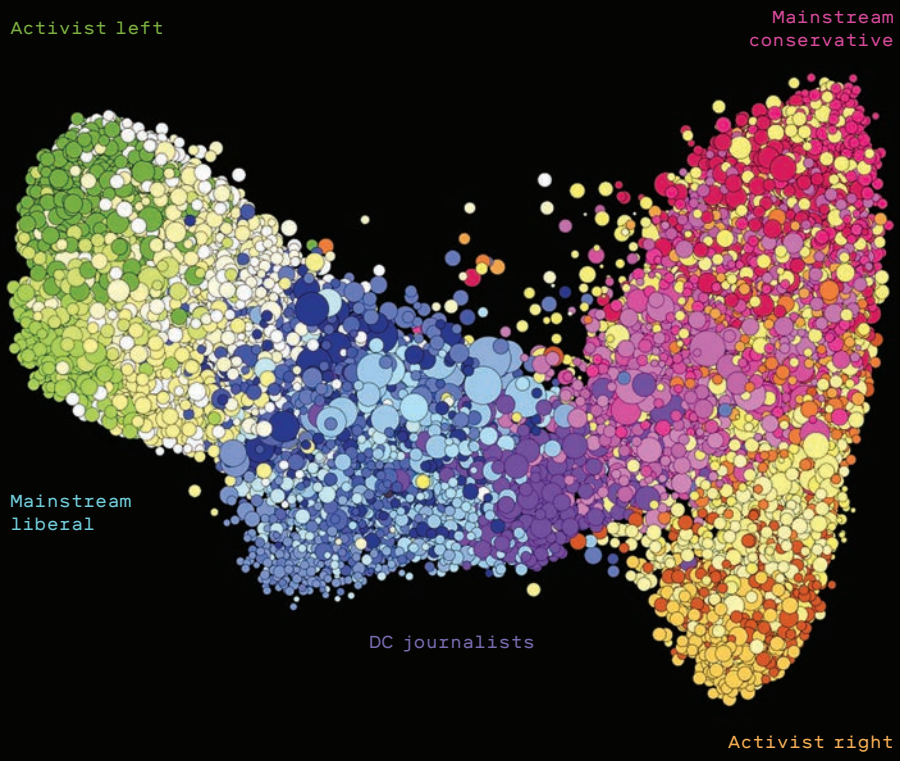

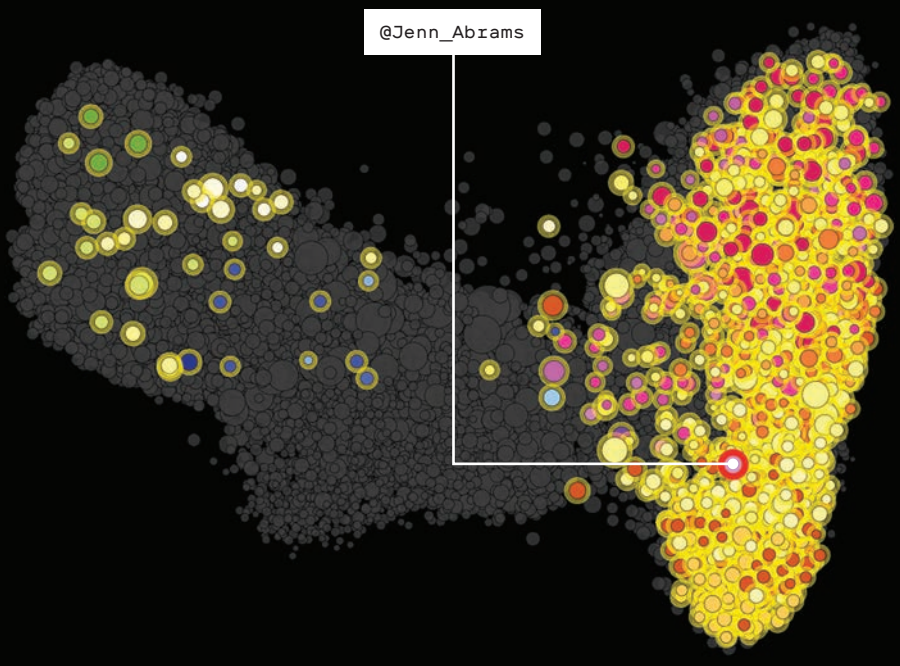

The Iranian blogosphere

The echo-chamber effect on the internet is nothing new. This 2008 map depicts the blogosphere in Iran, clustering together blogs that link to each other and coloring them by their content. Before a sustained government crackdown on online speech, supporters (lower right) and detractors (lower left) of the clerical regime each enjoyed substantial followings.

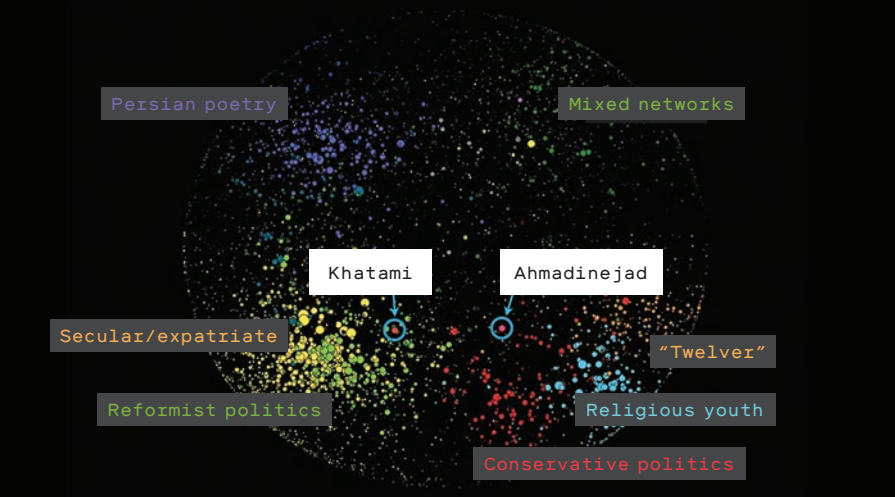

Turkish Twitter

This Twitter map of the political landscape in Turkey, analogous to the US map on pages 22-23, shows multidimensional polarization, with a dense sphere of influence around Erdogan supporters, on the far right side of the map, and two different poles of opposition on the other. These "amplification cores" of highly connected accounts have a disproportionate influence on the conversation and can rapidly boost polarizing messages.

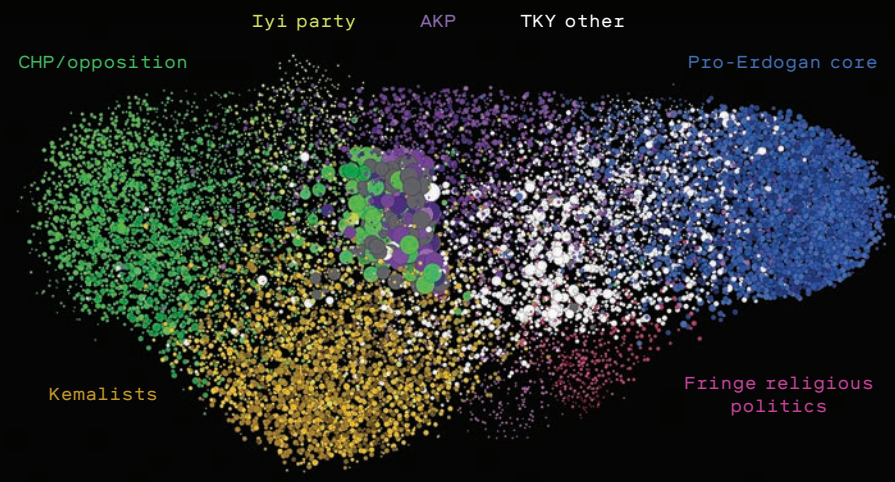

Russia: The same but different

This Twitter map of the Russian political landscape shows polarization in another context. There are clear pro- and anti-Putin clusters, but they're knit together by a broad set of mainly pro-government news and discussion-oriented accounts. A halo of apparently automated "personal" and marketing accounts surround the Putin fans.

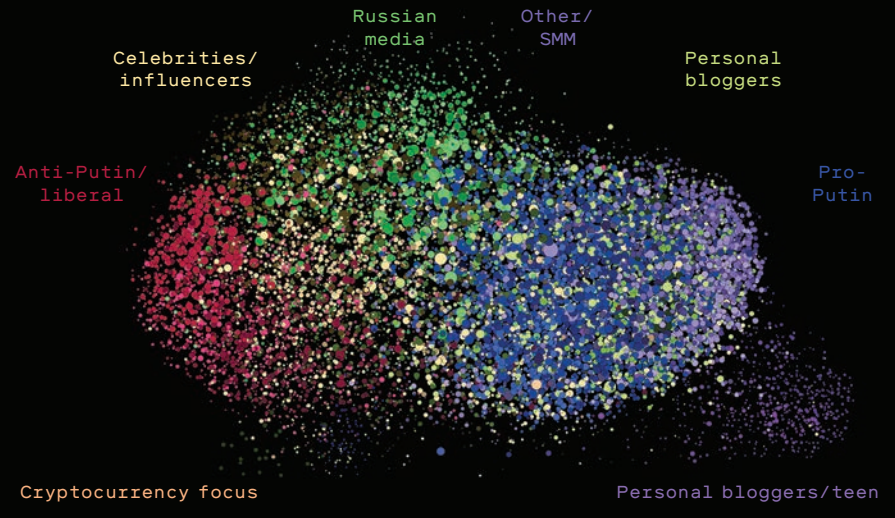

From 2008 to 2018

The first Obama campaign kicked off a technological revolution

What a difference a decade makes.

Consider: In 2008, the iPhone was less than a year old. BlackBerrys and e-mail dominated the palms of corporate and political information junkies. Television continued to be the dominant medium for political advertising and debates. Social media was a curiosity; governments and politicians who used it were still something of a novelty. It took the protests of the 2009 #IranElection—*Time* magazine dubbed Twitter "the medium of the movement"—to make mainstream journalists and politicians realize that smartphones and internet connections were fundamentally shifting how we lived, worked, played, advocated, campaigned, and governed.

Since then we've been living through probably the most rapid evolution of political campaigning in recent history. In each US election cycle, the technology used has advanced and morphed; the tools that gave Barack Obama the edge in 2008 and 2012 are very different from the ones that nudged Donald Trump to victory in 2016.

So where might things go next? What lessons will candidates for Congress in November's midterms have taken from Trump's victory? How will algorithm changes at a Facebook chastened by the Cambridge Analytica scandal alter the mechanics of influencing voters? Will 2020 or 2024 look as different from 2016 as 2016 did from 2008?

First, here's a brief history of a heady decade.

08

in electioneering. Where is it going next? By Alex Howard

20

2008

The key technological innovation that brought Barack Obama to the White House wasn't his tweets or a smartphone app. It was the Obama campaign's novel integration of e-mail, cell phones, and websites. The young, technology-savvy staffers didn't just use the web to convey the candidate's message; they also enabled supporters to connect and self-organize, pioneering the ways grassroots movements would adapt and adopt platforms in the campaign cycles to come.

The Obama campaign integrated social-networking features into My.BarackObama.com, where supporters could form groups, raise money, organize local events, and get information on the voters in their neighborhoods. The campaign converted online energy into offline action, from virtual phone bank rallies to voter registration to get-out-the-vote drives during primaries and the general election.

Of course, it didn't hurt that Obama was a unique and inspirational candidate to many—young, charismatic, and African-American. But it was adopting existing modern technology that helped an agile, insurgent campaign defeat Hillary Clinton, a member of one of America's political dynasties, in the primaries and then John McCain, a popular war hero, in the general election.

2012

The 2012 campaigns drew on these technologies still further. TV remained the dominant political medium for debates, but by August 2012, a majority of US adults were on Facebook. That meant more voters could now watch the social-media chatter accompanying the debates on a phone or computer screen, creating new opportunities for the campaigns to respond in real time and craft fund-raising appeals or amplify the messages that had resonated most.

Once again, the Obama campaign built a dream team of nerds to create the software that drove many aspects of the campaign. From messaging to fund-raising to canvassing to organizing to targeting resources to key districts and media buys, the reelection effort took the political application of data science to unprecedented heights. The Obama team created sophisticated analytic models that personalized social and e-mail messaging using data generated by social-media activity.

The Republican side, too, tried to create smarter tools, but it botched them. The Romney campaign's "Orca," a platform for marshaling volunteers to get out the vote on election day, suffered severe technical problems, becoming a cautionary tale of how not to manage a large IT project. For the moment, the technology gap between Democrats and Republicans remained wide.

2016

In many ways, Hillary Clinton's presidential campaign operation was a descendant of Obama's.

A big team of engineers led by ex-Googler Stephanie Hannon built dozens of tools, with a special focus on voter registration and turnout, and similarly formed an analytics unit to inform campaign decisions. The Clinton team built upon the institutional knowledge of the Democratic Party, trying to optimize and improve little things instead of developing a new "killer app."

In contrast, for all of Trump's prowess at messaging on Twitter, his campaign was an improvisational, bare-bones operation. Whereas Obama's and Clinton's teams poured resources into building their own systems in 2012 and 2016, Trump's campaign chose off-the-shelf tools and everyday vendors. It used social-media platforms and relatively simple websites to target voters, with data acquired from Facebook apps and targeting tools designed for commercial advertisers.

Just how much Cambridge Analytica helped in this effort is still highly contested. The firm, which worked for the Trump campaign, boasted that its "psychographic profiles," assembled using data that turned out to have been purloined from Facebook by an academic, contained as many as 5,000 data points on each of 220 million Americans. Yet Brad Parscale, who ran Trump's digital operation and has been named his 2020 campaign manager, has repeatedly insisted the campaign didn't use those profiles, relying instead on data from the Republican Party.

It's also hard to judge the impact of the $100,000 or so in "dark ads" that Facebook confirmed Russia ran on its platform in 2016. Their cost was tiny compared with the tens of millions spent on Facebook by the Clinton and Trump campaigns, and vanishingly small alongside the $6.5 billion that OpenSecrets estimates as the cost of the entire 2016 election cycle.

While the Trump campaign didn't have dozens of engineers and analysts on staff, however, it had something else that helped to close a yawning technological gap. These were the "embeds"—employees from Facebook, Twitter, and Google, picked for their Republican sympathies, who worked directly in the campaign's offices, teaching staffers how to get the most out of the platforms. The Clinton campaign was offered embeds but chose not to accept them. While it also spent tens of millions of dollars advertising on Facebook, it was much less sophisticated about it.

We have Facebook's own word for this. An internal company white paper that Bloomberg News obtained earlier this year reported that from June to November of 2016, Clinton's campaign tested 66,000 distinct ads while Trump's tested 5.9 million. Parscale told CBS's *60 Minutes* that they tested, on average, 50,000 to 60,000 different ads *daily*. This, in the words of the memo, "better leveraged Facebook's ability to optimize for outcomes."

On top of this, of course, the Trump campaign had Trump himself, whose personal communication style turned out to be perfectly suited both to social media and to the political moment. His capacity for regularly provoking outrage won him $5.9 billion worth of free attention from the mainstream media over the whole campaign, more than twice as much as Clinton, according to the analytics firm mediaQuant.

And it would be a mistake to forget that media attention, and the news that drives it, still matter more than any Facebook ad campaign. The scandals over Hillary Clinton's private e-mail server and the leaks of e-mails from the Democratic National Committee and her campaign chairman—believed to be the work of Russian hackers—generated vast amounts of coverage at critical moments, possibly enough to have swung the vote in Trump's favor.

2018

This year, campaigns will again deploy a broad range of technological tools to find and communicate with voters, using what they know about them to personalize ads, calls to volunteer and vote, and requests to donate. But the targeting is becoming ever more effective as more data on voters becomes available and tools for using it get better and more numerous.

The Democratic National Committee has set up a marketplace of vendors for the party's candidates. Left-leaning Higher Ground Labs has a portfolio of incubated startups offering polling, cheaper advertising, "persuasion science," and fund-raising tools like CallTime.ai, which tries to apply artificial intelligence to attracting donations. On the right-leaning Lincoln Network's app marketplace, many of the same vendors, from Salesforce to the community-management platform Nationbuilder, are available to conservative campaigns.

While the main factor in the midterms will be the voters' judgments on the Trump administration, smart uses of technology could make a difference in tight races. The tools available allow political novices to quickly gain direct access to attention and fund-raising if their candidacies and messages resonate with people—even if they are ignored by media outlets. Alexandria Ocasio-Cortez, a 28-year-old community activist, unseated a 20-year incumbent who outspent her fivefold in New York's Democratic primary for Congress, thanks in part to her viral video.

After the Russian "dark ads" and Cambridge Analytica scandals, will online campaigning at least be less murky? Yes, but not by much. Twitter and Facebook have made political ads more transparent, enabling the public to see who has bought them and putting advertisers through a vetting process; Google is expected to follow suit. But true transparency would mean having a file of all paid political ads, on a public website, with bulk open data downloads and an application programming interface (API) so that people could get the data without having to log into Twitter or Facebook themselves. It would also be backed up by a law instead of being voluntary. (Disclosure: I have a position on this, since while at the Sunlight Foundation I helped senators draft the Honest Ads Act. If enacted, the bill would not only mandate disclosures and disclaimers but update the definition of electioneering to include online platforms.)

In any case, these moves address only a tiny part of the problem. Political organizations and foreign states have long been able to channel "dark money" into political campaigns through nonprofits without identifying the source, and in July of this year, the US Treasury relaxed those rules even further. Regulating ads on social media also wouldn't address disinformation by foreign states or reverse the various Supreme Court decisions that have weakened US campaign-finance laws over the past decade.

2020 and beyond

Expect the campaigns in the next presidential race to use not radically new sorts of tools but more of the same: more data, better algorithms, and, consequently, more fine-grained targeting of voters, especially those judged to be crucial to tipping a district or a state in a candidate's favor.

What will probably evolve faster are the ways messages to those voters are created and spread. There may be gimmicks such as virtual-reality town-hall meetings or geotargeting—sending voters ads on their phones when they're near polling places or campaign events, for example. But the technology that's likely to have the most impact is something seemingly less advanced: video.

Many more people now have good enough mobile broadband to stream high-quality video on their phones than did just a few years ago. That's part of why unknown candidates on small budgets, such as Ocasio-Cortez, can become overnight sensations. As mobile video becomes more popular, though, it's also going to be exploited more as a tool of misinformation. Readily available software for creating video "deepfakes," such as the head of one person digitally swapped onto the body of another, is rapidly improving (see "Fake America great again," page 36). Generative adversarial networks (GANs), AI tools that pit two algorithms against each other, can be used to automate the creation of entirely artificial but believable imagery from scratch.

If the 2016 presidential race brought "fake news" into the lexicon, in 2020 the struggle to distinguish it from reality will reach a new level. For companies like Facebook, already under siege for permitting conspiracies and hate speech to circulate on their platforms, this may finally force a reckoning with society—and with legislators and regulators—about their responsibility, as the world's largest purveyors of information, to prevent the spread of personalized disinformation.

Alex Howard is a writer and open-government advocate based in Washington, DC, and former deputy director of the Sunlight Foundation.

The events of 2016 shattered the worldview of the internet idealists. Now they are casting around for alternatives.

By **Tim Hwang**

A field guide to the depressed former internet optimists

A long time ago, in the bad old days of the 2000s, debates about the internet were dominated by two great tribes: the Optimists and the Pessimists.

"The internet is inherently democratizing," argued the Optimists. "It empowers individuals and self-organizing communities against a moribund establishment."

"Wrong!" shouted the Pessimists. "The internet facilitates surveillance and control. It serves to empower only governments, giant corporations, and on occasion an unruly, destructive mob."

These battles went on at length and were invariably inconclusive.

Nevertheless, the events of 2016 seem to have finally shattered the Optimist consensus. Long-standing concerns about the internet, from its ineffectual protections against harassment to the anonymity in which teenage trolls and Russian spies alike can cloak themselves, came into stark relief against the backdrop of the US presidential election. Even boosters now seem to implicitly accept the assumption (accurate or not) that the internet is the root of multiple woes, from increasing political polarization to the mass diffusion of misinformation.

Tim Hwang is director of the Ethics and Governance of AI Initiative, a joint program of the MIT Media Lab and Harvard's Berkman Klein Center. (He is not to be confused with Tim Hwang, CEO of FiscalNote, profiled on page 56.)

All this has given rise to a new breed: the Depressed Former Internet Optimist (DFIO). Everything from public apologies by figures in the technology industry to informal chatter in conference hallways suggests it's become very hard to find an internet Optimist in the old, classic vein. There are now only Optimists-in-retreat, Optimists-in-doubt, or Optimists-hedging-their-bets.

As Yuri Slezkine argues wonderfully in *The House of Government*, there is a process that happens among believers everywhere, from Christian sects to the elites of the Russian Revolution, when a vision is unexpectedly deferred. Ideologues are forced to advance a theory to explain why the events they prophesied have failed to come to pass, and to justify a continued belief in the possibility of something better.

Among the DFIOs, this process is giving rise to a boomlet of distinct cliques with distinct views about how the internet went wrong and what to do about it. As an anxiety-ridden DFIO myself, I've been morbidly cataloguing these strains of thinking and have identified four main groups: the Purists, the Disillusioned, the Hopeful, and the Revisionists.

These are not mutually exclusive positions, and most DFIOs I know combine elements from them all. I, for instance, would call myself a Hopeful-Revisionist.

The question is, do these tribes matter? Or are the Pessimists ultimately right that the internet is essentially destructive to society? Does the flowering of DFIO cliques, as Slezkine's book suggests, simply represent the final throes of a dying movement?

I say no. Both Optimism and Pessimism make the mistake of assuming that the internet has inherent features, but like any technology conceived of and built by humans, it is shaped by human struggles, by the push and pull of a multitude of interests and schools of thought. What's needed is a coalition around a New Optimism—one that celebrates what's working, is honest about what isn't, and articulates a path forward grounded not so much in technological fixes as in a richer understanding of trust, identity, and community.

1 The Purists

"The internet was a wonderful place before it became corrupted by corporations/commercialization/etc." You hear this trope frequently among some DFIOs. Purists are still true believers—they think the "heart" of the technology, however defined, is something fundamentally good. The blame, in this view, lies with intervening forces that subverted the technology and prevented it from reaching its full promise. Purists want to launch the next great crusade and frequently talk about using blockchain for everything, breaking up the big tech firms, or putting an end to the scourge of advertising.

2 The Disillusioned

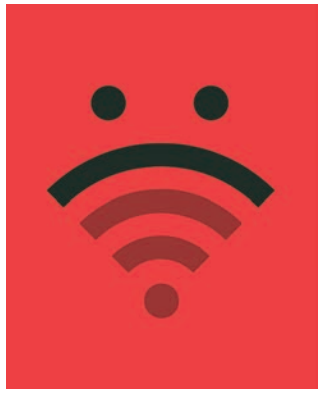

"The internet was never all that great," you'll sometimes hear one DFIO say to another. "We're realizing that now." While the Purists maintain that a golden era of the internet really did exist, the Disillusioned unhappily believe that there was never any substance to these claims. Close cousins include the Saw-It-All-Alongers, former Optimists who also want the feel-good kicks of saying that everyone else is catching up to what they divined years ago. You'll frequently find members of both groups enthusiastically using social media to hate on social media.

3 The Hopeful

One response to a perceived local failure is to seek optimism globally. This is the aspiration of the Hopeful, who try to vindicate the dreams of internet Optimism by foraging for positive moments in the wider world of the web. Some will point to the mass departure of younger generations from platforms like Facebook, or to intriguing experiments in digital democracy in other countries, or to the vitality and diversity of internet culture in general, as signs that a brighter day is yet to come. The Hopeful love Tumblr unconditionally, shared the lemon-rolling video nostalgically, and collect whimsical Slack memberships like they're going out of style.

4 The Revisionists

Many Optimists believed that the structure of the internet by itself—manifested in collaborative projects such as wikis or crowdfunding—would bend social outcomes in their favor. One response to the events of 2016 has been to revisit this assumption, claiming that while the basics might have been right, more work is needed to realize the original vision. Revisionists want to preserve the original aspirations for the web through amendment, calling for a new effort to design better communities and systems for governing society online. They extol the virtues of stronger community guidelines, ways to influence behavior through "nudging" interfaces, and the power of user-centered design.

Wharton
University of Pennsylvania
Aresty Institute of Executive Education

EXECUTIVE EDUCATION

excel

verb \ik·'sel\

"The moment I realized I had become part of something that would take me further than I've ever been."

Define *your* Wharton moment.

Wharton's ***General Management Program*** provides a flexible learning journey for distinguished senior executives ready for greater challenges. You will receive expert one-on-one **executive coaching** in Wharton's rigorous academic environment and return to your organization with an enriched global perspective and in-depth strategies for immediate success. Upon successful completion of the program, you will be awarded **Wharton alumni status** and join a powerful network of 96,000 peers in over 150 countries.

EXCEL AT A HIGHER LEVEL:

EXECED.WHARTON.UPENN.EDU/**GMP**

General Management Program

A FLEXIBLE LEARNING JOURNEY

Design your own curriculum of 6 programs within 2 years. Choose from 30+ eligible programs in:

- LEADERSHIP
- FINANCE
- STRATEGY & INNOVATION
- MARKETING

Where we are now

Deepfakes are so easy, even we made one (p36). When more surveillance might mean (slightly) more freedom (p42). A list of the holes in America's voting systems (p44). Kenya's lesson for political tech solutionists (p48). China's quest to create the perfectly measurable, governable society (p50). And the whiz kid using data to upend DC lobbying (p56).

2


FAKE AMERICA
GREAT AGAIN

Inside the race to catch the worryingly real fakes that can be made using artificial intelligence.

BY WILL KNIGHT
PHOTOGRAPH BY BRUCE PETERSON

Guess what? I just got hold of some embarrassing video footage of Texas senator Ted Cruz singing and gyrating to Tina Turner. His political enemies will have great fun showing it during the midterms. Donald Trump will call him "Dancin' Ted."

Okay, I'll admit it—I created the video myself. But here's the troubling thing: making it required very little video-editing skill. I downloaded and configured software that uses machine learning to perform a convincing digital face-swap. The resulting video, known as a deepfake, shows Cruz's distinctively droopy eyes stitched onto the features of actor Paul Rudd doing lip-sync karaoke. It isn't perfect—there's something a little off—but it might fool some people (watch it at technologyreview.com/cruzrudd).

Photo fakery is far from new, but artificial intelligence will completely change the game. Until recently only a big-budget movie studio could carry out a video face-swap, and it would probably have cost millions of dollars. AI now makes it possible for anyone with a decent computer and a few hours to spare to do the same thing. Further machine-learning advances will make even more complex deception possible—and make fakery harder to spot.

These advances threaten to further blur the line between truth and fiction in politics. Already the internet accelerates and reinforces the dissemination of disinformation through fake social-media accounts. "Alternative facts" and conspiracy theories are common and widely believed. Fake news stories, aside from their possible influence on the last US presidential election, have sparked ethnic violence

These still images of Ted Cruz and Paul Rudd are taken from the footage that was fed to a face-swapping program.

in Myanmar and Sri Lanka over the past year. Now imagine throwing new kinds of real-looking fake videos into the mix: politicians mouthing nonsense or ethnic insults, or getting caught behaving inappropriately on video—except it never really happened.

"Deepfakes have the potential to derail political discourse," says Charles Seife, a professor at New York University and the author of *Virtual Unreality: Just Because the Internet Told You, How Do You Know It's True?* Seife confesses to astonishment at how quickly things have progressed since his book was published, in 2014. "Technology is altering our perception of reality at an alarming rate," he says.

Are we about to enter an era when we can't trust anything, even authentic-looking videos that seem to capture real "news"? How do we decide what is credible? Whom do we trust?

Real fake

Several technologies have converged to make fakery easier, and they're readily accessible: smartphones let anyone capture video footage, and powerful computer graphics tools have become much cheaper. Add artificial-intelligence software, which allows things to be distorted, remixed, and synthesized in mind-bending new ways. AI isn't just a better version of Photoshop or iMovie. It lets a computer learn how the world looks and sounds so it can conjure up convincing simulacra.

I created the clip of Cruz using OpenFaceSwap, one of several face-switching programs that you can download for free. You need a computer with an advanced graphics chip, and this can set you back a few thousand bucks. But you can also rent access to a virtual machine for a few cents per minute using a cloud machine-learning platform like Paperspace. Then you simply feed in two video clips and sit back for a few hours as an algorithm figures out how each face looks and moves so that it can map one onto the other. Getting things to work is a bit of an art: if you choose clips that are too different, the result can be a nightmarish mishmash of noses, ears, and chins.

But the process is easy enough.

Face-swapping was, predictably, first adopted for making porn. In 2017, an anonymous Reddit user known as Deepfakes used machine learning to swap famous actresses' faces into scenes featuring adult-movie stars, and then posted the results to a subreddit dedicated to leaked celebrity porn. Another Reddit user then released an easy-to-use interface, which led

Getting things to work is a bit of an art: if you choose clips that are too different, the results can be a mishmash of noses, ears, and chins.

CNN / THE TONIGHT SHOW WITH JIMMY FALLON

to a proliferation of deepfake porn as well as, for some odd reason, endless clips of the actor Nicolas Cage in movies he wasn't really in. Even Reddit, a notoriously freewheeling hangout, banned such nonconsensual pornography. But the phenomenon persists in the darker corners of the internet.

OpenFaceSwap uses an artificial neural network, by now the go-to tool in AI. Very large, or "deep," neural networks that are fed enormous amounts of training data can do all sorts of useful things, including finding a person's face among millions of images. They can also be used to manipulate and synthesize images.

OpenFaceSwap trains a deep network to "encode" a face (a process similar to data compression), thereby creating a representation that can be decoded to reconstruct the full face. The trick is to feed the encoded data for one face into the decoder for the other. The neural network will then conjure, often with surprising accuracy, one face mimicking the other's expressions and movements. The resulting video can seem wonky, but OpenFaceSwap will automatically blur the edges and adjust the coloring of the newly transplanted face to make things look more genuine.

Similar technology can be used to re-create someone's voice, too. A startup called Lyrebird has posted convincing demos of Barack Obama and Donald Trump saying entirely made-up things. Lyrebird says that in the future it will limit its voice duplications to people who have given their permission—but surely not everyone will be so scrupulous.

There are well-established methods for identifying doctored images and video. One option is to search the web for images that might have been mashed together. A more technical solution is to look for telltale changes to a digital file, or to the pixels in an image or a video frame. An expert can search for visual inconsistencies—a shadow that shouldn't be there, or an object that's the wrong size.

Dartmouth University's Hany Farid, one of the world's foremost experts, has shown how a scene can be reconstructed in 3-D in order to discover physical oddities. He has also proved that subtle changes in pixel intensity in a video, indicating a person's pulse rate, can be used to spot the difference between a real person and a computer-generated one. Recently one of Farid's former students, now a professor at the State University of New York at Albany, has shown that irregular eye blinking can give away a face that's been manipulated by AI.

Still, most people can't do this kind of detective work and don't have time to study every image or clip that pops up on Facebook. So as visual fakery

A startup has posted convincing demos of Barack Obama and Donald Trump saying entirely made-up things.

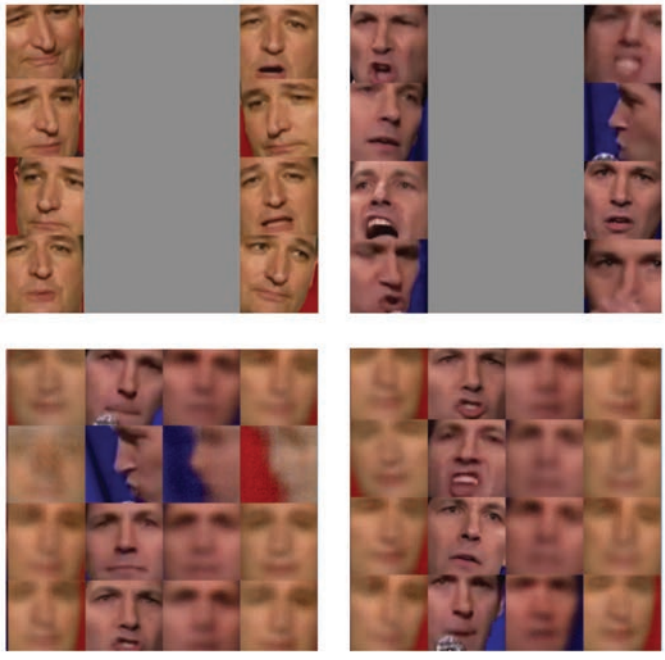

OpenFaceSwap previews attempted face swaps during training. Early tries can often be a bit weird and grotesque.

GANs can turn daytime scenes into nighttime ones and dream up imaginary celebrity faces.

The software takes several hours to produce a good face swap. The more training data, the better the end result.

has become more common, there's been a push to automate the analysis. And it turns out that not only does deep learning excel at making stuff up, it's ideal for scrutinizing images and videos for signs of fakery. This effort has only just begun, though, and it may ultimately be hindered by how realistic the automated fakes could become.

Networks of deception

One of the latest ideas in AI research involves turning neural networks against themselves in order to produce even more realistic fakes. A "generative adversarial network," or GAN, uses two deep neural networks: one that's been trained to identify real images or video, and another that learns over time how to outwit its counterpart. GANs can be trained to produce surprisingly realistic fake imagery.

Beyond copying and swapping faces, GANs may make it possible to synthesize entire scenes and people that look quite real, turning a daytime scene into a nighttime one and dreaming up imaginary celebrities. GANs don't work perfectly, but they are getting better all the time, and this is a hot area of research (*MIT Technology Review* named GANs one of its "10 Breakthrough Technologies" for 2018).

Most worrying, the technique could also be used to evade digital forensics. The US's Defense Advanced Research Projects Agency invited researchers to take part in a contest this summer in which some developed fake videos using GANs and others tried to detect them. "GANs are a particular challenge to us in the forensics community because they can be turned against our forensic techniques," says Farid. "It remains to be seen which side will prevail."

This is the end

If we aren't careful, this might result in the end of the world—or least what seems like it.

In April, a supposed BBC news report announced the opening salvos of a nuclear conflict between Russia and NATO. The clip, which began circulating on the messaging platform WhatsApp, showed footage of missiles blasting off as a newscaster told viewers that the German city of Mainz had been destroyed along with parts of Frankfurt.

It was, of course, entirely fake, and the BBC rushed to denounce it. The video wasn't generated using AI, but it showed the power of fake video, and how it can

spread rumors at warp speed. The proliferation of AI programs will make such videos far easier to make, and even more convincing.

Even if we aren't fooled by fake news, it might have dire consequences for political debate. Just as we are now accustomed to questioning whether a photograph might have been Photoshopped, AI-generated fakes could make us more suspicious about events we see shared online. And this could contribute to the further erosion of rational political debate.

In *The Death of Truth*, published this year, the literary critic Michiko Kakutani argues that alternative facts, fake news, and the general craziness of modern politics represent the culmination of cultural currents that stretch back decades. Kakutani sees hyperreal AI fakes as just the latest heavy blow to the concept of objective reality.

"Before the technology even gets good, the fact that it exists and is a way to erode confidence in legitimate material is deeply problematic," says Renee DiResta, a researcher at Data for Democracy and one of the first people to identify the phenomenon of politically motivated Twitter misinformation campaigns.

Perhaps the greatest risk with this new technology, then, is not that it will be misused by state hackers, political saboteurs, or Anonymous, but that it will further undermine truth and objectivity itself. If you can't tell a fake from reality, then it becomes easy to question the authenticity of anything. This already serves as a way for politicians to evade accountability.

President Trump has turned the idea of fake news upside down by using the term to attack any media reports that criticize his administration. He has also suggested that an incriminating clip of him denigrating women, released during the 2016 campaign, might have been digitally forged. This April, the Russian government accused Britain of faking video evidence of a chemical attack in Syria to justify proposed military action. Neither accusation was true, but the possibility of sophisticated fakery is increasingly diminishing the credibility of real information. In Myanmar and Russia new legislation seeks to prohibit fake news, but in both cases the laws may simply serve as a way to crack down on criticism of the government.

As the powerful become increasingly aware of AI fakery, it will become easy to dismiss even clear-cut video evidence of wrongdoing as nothing more than GAN-made digital deception.

The truth will still be out there. But will you know it when you see it?

Will Knight is a senior editor at MIT Technology Review who covers artificial intelligence.

The final video shows Ted Cruz's face stitched almost seamlessly onto Paul Rudd (technologyreview.com/cruzrudd).

Perhaps the greatest risk is that the technology will further undermine truth and objectivity.

China's adoption of ever more intrusive technology could, paradoxically, lead to stronger civil liberties.

By **Yasheng Huang**

Can big data tame Big Brother?

By 2020, China's new system of social credit scoring is expected to give each citizen a trustworthiness rating based on anything from shopping habits to choice of friends. It may seem like an ideal tool for an authoritarian government that wants to control its citizens. But while authoritarian regimes have always been enthusiastic adopters of surveillance technology, in China's case big data may (inadvertently) make the country a little less repressive.

Yasheng Huang is a professor in international management at the MIT Sloan School of Management. He is also the author of Capitalism with Chinese Characteristics: Entrepreneurship and the State.

Privacy is now a thing in China

A few years ago, in an article for the *Boston Globe* called "How Privacy Became an American Value," historian and author Ted Widmer detailed how Americans inherited and amplified the British sense of privacy—the idea of "keeping oneself to oneself." Things such as income, health, and leisure pursuits are routinely considered private—most of all from the government. The Fourth Amendment to the US Constitution prohibits "unreasonable searches and seizures," promising that Americans have the right "to be secure in their persons, houses, papers, and effects."

Although it's dangerous to make generalizations, it's fair to say Chinese and Western citizens differ in the weights they assign to privacy, individual rights, and freedom of speech. It's not that Chinese people don't value those things; it's just that they may value things like economic growth and income more. (In the West, too, people are sometimes willing to trade rights for other benefits. In a survey earlier this year by Credible, a personal-finance website, nearly half of American millennials said they would give up their vote in the next two presidential elections in exchange for having their student loans forgiven.)

One reason Chinese attitudes are different is that as recently as the 1980s, the word "privacy" had negative connotations in China. Chinese norms are anchored in 2,000 years of a Confucian culture that values the intensity of interpersonal relationships. One way to solidify those relationships is through transparency and full disclosure. A circumstance that triggers secrecy is typically an unsavory one. If something is good, why not tell us? Privacy in this context was equated with preserving a dirty secret. To be private was to be antisocial.

Yet as social interactions have evolved in China, so have Chinese values. The rise of big-data technology in China has contributed to a far more acute awareness of privacy than other—momentous—socioeconomic developments such as GDP growth, globalization, and urbanization.

The reason for this is that big data has decisively broken the personal intimacy of Confucian culture. On WeChat, you can friend thousands of people you barely know. On Alibaba, you can do business with people you wouldn't recognize if they knocked on your door. The digital economy is impersonal to an unprecedented degree, and as a consequence the old Confucian social contract, built on solidifying personal relationships by telling your neighbors everything, has crumbled. Though it may threaten privacy, big data has also brought unprecedented attention to the very notion of privacy. In the long run, this may be the force that undermines Big Brother.

China is a repressive society with or without big data. Technology has made the repression more precise, but that might be an improvement over indiscriminate repression.

Better or worse than what?

China's surveillance culture existed long before the rise of big data. In his book *The Government Next Door*, Luigi Tomba details how Chinese politics have been micromanaged at the neighborhood level. Residential communities are monitored by neighborhood committees performing semigovernmental functions: reporting dissent, resolving conflicts, and managing both petitions to the government and protests against it. These functions used to be the task of retired elderly women, whom the former *Wall Street Journal* reporter Adi Ignatius memorably called the "small-feet KGB." (In traditional China, women had their feet bound at birth.) The question is whether monitoring and repression through impersonal technology is better or worse than these personal intrusions.

One of the most important roles of the small-feet KGB was to enforce China's one-child policy. The Chinese fertility rate fell dramatically while the policy applied, from 1979 to 2015—a testament to the effectiveness of these personal surveillance tactics.

In ancient China, there was a joint liability system under which three to five households were linked together. If a member of one household committed an offense, all the households were punished. During the Cultural Revolution, punishments for political dissenters were routinely meted out to their immediate family members. The political system compensated for a lack of data on individual activities by deterring dissent broadly and harshly.

Big data would be a threat if Chinese citizens could be expected to have an abundance of political and civil liberties in its absence. But China is a repressive, authoritarian society with or without big data. Technology has made the repression more precise, but precise repression might be an improvement over indiscriminate repression.

Traffic is expensive

In a segment on *The Late Show* earlier this year, comedian Stephen Colbert told his audience that the social credit scores being rolled out in China would dock citizens for instances of jaywalking, among other things. That might sound harsh, but then Colbert has evidently never driven in Beijing.

Alibaba, China's largest online retailer, is using cloud computing to combat China's suffocating traffic. In 2016, the company introduced a traffic management system called City Brain in Hangzhou, where Alibaba is headquartered. Unlike Google Maps, City Brain is a collaborative project with the city government; it can tap into the traffic and transportation bureaus' systems for video footage of traffic incidents. The municipal government relies on City Brain to identify the best routes for emergency vehicles and to plan new roads and bus routes.

Might City Brain also be used for some Big Brother–ish functions? Probably, but easing China's traffic nightmares and getting emergency patients to the hospital faster aren't trivial gains. According to the Chinese transportation ministry, traffic congestion in 2017 cost about 20 percent of total urban disposable income, or about 5 to 7 percent of China's GDP. About 20 percent of the gasoline consumed in China is wasted.

The social benefits gained through big-data technology don't obviate the political downsides. The question is: just how "down" is the downside, and how "up" is the upside?

Your vote is in jeopardy

Cyberattacks on the 2016 US election caused states to bolster the defenses of their voting systems. It hasn't been enough.

By Martin Giles
Portrait by Lyndon French

Russian hackers targeted US electoral systems during the 2016 presidential election. Much has been done since then to bolster those systems, but J. Alex Halderman, director of the University of Michigan's Center for Computer Security and Society, says they are still worryingly vulnerable (see "How hackers could cause chaos," page 46). *MIT Technology Review*'s Martin Giles discussed election security with Halderman, who has testified about it before Congress and evaluated voting systems in the US, Estonia, India, and elsewhere.

Lots of things, from gerrymandering to voter ID disputes, could undermine the integrity of the US electoral process. How big an issue is hacking in comparison?
Things like gerrymandering are a question of political squabbling within the rules of the game for American democracy. When it comes to election hacking, we're talking about attacks on the United States by hostile foreign governments. That's not playing by the rules of American politics; that's an attempt to subvert the foundations of our democracy.

How much has election security improved since the 2016 US presidential election?
One thing that's improved is awareness. States are taking the first necessary steps to protect their systems—things like making sure they run vulnerability scans on software, and that electoral staff have security clearance to receive threat intelligence from the federal government. Progress accelerated in March when Congress allocated $380 million in new funding that will help states afford to upgrade insecure equipment and make other improvements, but there's still a lot more work to be done.

What element of the voting process worries you the most?
The part that keeps me up at night is the electronic voting machines. Every machine has to be programmed with the ballot design, and that programming is copied in by election officials on a USB stick or memory card. If someone can infect that programming, they can spread an attack to the machines and potentially tamper with a fraction of the votes without anyone detecting it.

So what can be done to address this risk?
We need to make sure that every vote is recorded on a piece of paper, too. Without paper, there may be no evidence we can go back and look at that would reveal vote tampering. We also need to make attacks as difficult as possible by making sure systems used to program ballot design are locked down and never accessible from the internet.

What other areas beyond voting machines are vulnerable?
Voter registration systems connected to the internet are a major concern. In 2016, one of the most worrying cyberattacks was Russian attempts to probe, and in some cases hack into, voter registration databases. We also need to worry about electronic poll books that many states use to check voters in on Election Day. This equipment is often networked, and if it fails it could lead to chaos at the polls.

How can we bolster defenses here?
The main thing is to apply the same good security practices developed for protecting other government and industry databases. We also need to have backup procedures in place in case the technology fails.

Auditing results can catch vote manipulation. Are post-election audits in the US sufficiently robust?
No. Some states don't check ballots at all; others examine them in a fixed fraction of precincts, but in a close contest, that might not catch vote tampering concentrated in precincts that aren't checked. We need "risk-limiting" audits. Here you agree in advance the probability you're willing to tolerate of an election outcome being manipulated and not detected. You then look at enough paper ballots so the odds of someone getting away with fraud are lower than the target percentage.

Why don't we have these audits everywhere?
States have been slow to adopt new ways of countering cyberthreats. Fortunately, risk-limiting audits don't have to be particularly expensive. When an election isn't close, you might be able to confirm the result with high statistical confidence by examining a few hundred ballots across a state; in extremely

close elections, you often have to do an automatic recount anyway.

Would it be better if the US had a federally mandated, nationwide voting system rather than many different state and local ones?
It might be easier to secure a single, unified voting system, but election administration in the US is the responsibility of state and local governments, and I don't see that changing soon. What we can do is to set national standards for election cybersecurity that states should meet or exceed.

Could one tie federal money for securing elections to the adoption of those standards at the state level?
That could be quite effective, and there's a bipartisan draft bill in Congress called the Secure Elections Act that would do just that.

What would have to happen for online voting, Estonia-style, to become broadly viable in the US?
Online voting carries extremely big risks. You need to protect internet-connected servers running the election from sophisticated adversaries and protect voters' own devices from malware. That's why Estonia is the only country where national elections are largely online, and its system is unlikely to withstand a concerted attack. It may be decades before we're able to secure online systems to the same level we expect from voting in polling places today.

Some people have floated the idea of blockchain-based voting systems. Are you a fan?
Blockchain doesn't fix the hard parts of securing online elections. It's just another form of recording votes. If attackers compromise voters' devices or the servers that record votes and log them to the blockchain, they can still manipulate election outcomes. There are no easy solutions here.

Here's how hackers could cause chaos in the US midterm elections

In the months leading up to November's midterm elections in the US, hordes of foreign hackers will head to their keyboards in a bid to influence the outcome. Their efforts will include trying to get inside the digital infrastructure that supports the electoral process.

There's a worrying precedent here. Last year, the Department of Homeland Security (DHS) notified 21 states that Russia had targeted their election systems in the months leading up to the 2016 presidential election.

DHS officials said the Russians were mainly scanning computers and networks for security holes rather than taking advantage of any flaws they discovered. Still, that's no cause for complacency. Intelligence officials are already warning that Russia is intent on meddling in this year's midterm elections, too—and most of the digital technology that will be used predates the launch of the first iPhone in 2007. Here's what cyberattackers might target.

Voter registration systems

THE TECHNOLOGY: These systems keep a digital record of authorized voters, and data from them populates "poll books" used to check people in at precinct polling stations.

THE RISKS: Many voter registration systems are old: a report last year by the Brennan Center for Justice at New York University School of Law estimated that 41 states were still using ones built at least a decade ago. They are hosted on servers and need connectivity to receive voter data and transfer it to poll books. Hackers who gain access to them could erase voters' entries or create fictitious ones and then mail in votes for the fake personas. That could tip the balance in tight races.

This makes the systems tempting targets. In his indictment of 12 Russian hackers in July, US special counsel Robert Mueller alleged that they penetrated the website of one (unnamed) state board of elections in 2016 and stole partial Social Security numbers, driver's license numbers, and other data for around half a million voters.

Voter check-in

THE TECHNOLOGY: In many states, precinct poll workers use tablet-like electronic poll books, rather than paper ones, to verify voters. These machines are often networked to one another and run tailor-made software.

THE RISKS: Hackers could target the networks to gain access to poll books, either shutting them down or altering data that's on them. They could also break into the systems of companies that develop software for the poll books and insert malicious code.

Compromising poll books could cause chaos during an election. For instance, voters may be told that they've already voted when in fact they haven't. Ideally, all polling stations should have backup plans in place that allow them to print provisional ballots if the machines fail.

Voting machines

THE TECHNOLOGY: The US uses two main types of electronic voting machines. Optical-scan ballot readers scan and record paper ballots filled in by voters, while direct-recording electronic, or DRE, machines display ballot options on a screen and record voters' choices electronically. Only some DRE machines produce paper records too.

THE RISKS: Voting machines are programmed with the ballot design, which includes names of the races and candidates involved. The design is set up on election management systems at a central election office or a vendor. The information is typically then transferred to each machine by officials using memory cards or USB keys. Hackers can target the central computers to spread malicious code to multiple machines, or they can target individual devices.

If officials suspect optical-scan ballot readers have been hacked, they can check the paper ballots; with DREs, there's sometimes no paper record to look at. Paperless machines are still used in 13 states, and five rely solely on them.

Vote tallying and reporting

THE TECHNOLOGY: The software managing vote tallying and reporting typically runs on computers using standard operating systems.

THE RISKS: Hackers could target the software to throw doubt on the outcome of elections. While this may sound unlikely, there are strong suspicions Russian hackers were behind an attack that deleted key files from the Ukrainian central election commission's system in a 2014 vote.

The good news is that almost all US states check outcomes against reports from individual precincts before certifying official results. So any confusion sown by an attack on vote tallying and reporting software should eventually be resolved as long as the underlying voting processes remain secure.

Beyond all these risks, plenty of other nightmare scenarios could affect the different stages reviewed here. They include distributed denial of service attacks, which knock web-connected systems out of action by flooding them with fake traffic, and ransomware attacks, which use malware to encrypt data—or, in the worst case, destroy it.

SIMON LANDREIN

Long before the internet, hate speech flourished in echo chambers of a different kind.

By **Nanjala Nyabola**

Kenya's technology evolved. Its political problems stayed the same.

In 2007, incumbent Mwai Kibaki won a divisive presidential election in Kenya. Street protests escalated to ethnic violence in parts of the country, and by April 2008 more than 1,500 people had been killed. A decade later, another election also featured widespread allegations of fraud, and more violence. The casualties were lower this time, but just over 100 were killed, almost all by the police in opposition strongholds.

Technology and politics are inextricably linked in Kenya, in part because technology was proposed as the solution to the structural issues that led to the violence following the 2007 election. The Independent Review Commission established in 2008 argued that technology would help bridge the trust gap between key political actors and protect the autonomy of the bureaucracy surrounding elections. In line with the commission's recommendations, voter registration, voter identification, and vote tallying are all nominally computerized. These efforts culminated in the 2017 election, which was supposed to be Africa's first fully digital election.

That they didn't work is a lesson in what technology can and cannot fix in politics. Perhaps the primary lesson is that when communication platforms cater to specific audiences while shutting out others amid existing social divides, hate speech thrives. And this was as true of pre-internet technology as it is now.

Nanjala Nyabola is the author of Digital Democracy, Analogue Politics: How the Internet Era Is Transforming Kenya.

Twenty-five years ago, Kenya had one state-owned radio and television broadcaster, but by 2017 there were over 60 licensed television stations and 178 radio stations. Kenya has two official languages—English and Kiswahili—and at least 44 ethnic groups, most of which speak a third language that's often unintelligible to outsiders. As part of the democratization process in the 1990s, local-language radio stations were encouraged as way to bring more people into the national conversation while preserving regional cultures.

The unexpected result

But it turns out that local-language radio stations are particularly vulnerable to becoming channels for hate speech. They function as closed systems shielded from scrutiny by institutions that don't have the capacity to manage them, or any interest in doing so. By the time the threat inherent in these stations was identified, it had already materialized, and regulators are still playing catch-up. Kenya has laws prohibiting hate speech in media and numerous bodies ostensibly working against it, but as late as July 2017, the National Integration and Cohesion Commission was warning that Kameme FM, a Kikuyu-language radio station owned by Kenyan president Uhuru Kenyatta, was responsible for the most incidents of hate speech related to the election. Despite this finding, the station received no public censure.

Technology has changed radically in Kenya over the past decade, as it has everywhere else. Almost nine out of 10 people have a mobile phone, and a quarter of the homes have an Internet connection—among the highest rates in the developing world. In a population of about 48 million, there are at least seven million Kenyan Facebook accounts and another 10 million on WhatsApp. Twitter lags behind at only a million accounts, but

Technology reflects the values of the societies in which it's deployed, and can't fix problems that a society is unwilling to fix within itself.

it is supplanting television and print as the premier space for political critique. mPesa, the mobile money transfer platform, was less than six months old at the time of the 2007 election. Today, transactions on mPesa equal almost a third of the country's GDP, and Kenya has the highest number of mobile money transactions in the world.

A better way to manipulate

These technologies have transformed the way Kenyans produce and consume political information, and by extension the way they interact with each other. Yet the causes of hate speech on local-language radio stations haven't gone away. Many of the practices horrifying analysts in the West—the exploitation of identity politics to win elections, or the influence of money on decision-making in electoral politics, for example—are familiar to Kenyans. (Cambridge Analytica, which has been accused of using Facebook data to manipulate voters in the US and UK, has been present and active in Kenya since at least 2012.)

Technology has given powerful people a more effective way of influencing the electorate by using language that dehumanizes the "other," particularly on the basis of ethnicity, and then paying individuals to instigate violence that triggers reprisals. With the internet, information now travels faster, without the moderating impact of editing or verification. The same speed that makes social media the preferred avenue for alerting the public about emergencies makes it especially efficient at disseminating hateful opinions. Moreover, information on social media travels in relatively insular networks, which makes people less likely to encounter opinions that they disagree with or that challenge their thinking.

Division by design

According to experts like Zeynep Tufecki, some of the characteristics that encourage hate speech are embedded in the logic of these platforms, and indeed are what make them so attractive to advertisers; they are the features that make microadvertising and therefore profitability possible. Again, this did not begin with the internet; Kenya's radio market, similarly, was stratified to sell advertising.

Stratified audiences may be good for advertisers, but they weaken the role of media because public discourse among different groups is no longer starting from the same point. Stratification also makes these spaces harder for regulators to monitor. In 2007, Kenyans rioted in response to allegations of violence, broadcast on local-language radio, that turned out to be unfounded. In 2017, ethnically charged memes alleged that a genocide would happen in Kenya should certain candidates prevail. The absence of fact-checking, verification, and oversight on these platforms, coupled with the speed of transmission, allows inflammatory rumors to spread.

Kenya reminds us that with every evolution of technology, new concerns arise about its role in politics. Technology reflects the values of the societies in which it's deployed, and can't fix problems that a society is unwilling to fix within itself. The internet has sped up political discourse and further insulated it from scrutiny, mimicking and amplifying the experience Kenyans had with local-language radio stations, and reminding us that some problems are bigger than the medium.

Here's how China rules using data, AI, and internet surveillance.

Who needs democracy when you have data?

By
CHRISTINA LARSON

Photographs by
Gilles Sabrié

People in Beijing are always under the watchful eye of Mao—and myriad surveillance cameras.

In 1955, science fiction writer Isaac Asimov published a short story about an experiment in "electronic democracy," in which a single citizen, selected to represent an entire population, responded to questions generated by a computer named Multivac. The machine took this data and calculated the results of an election that therefore never needed to happen. Asimov's story was set in Bloomington, Indiana, but today an approximation of Multivac is being built in China.

For any authoritarian regime, "there is a basic problem for the center of figuring out what's going on at lower levels and across society," says Deborah Seligsohn, a political scientist and China expert at Villanova University in Philadelphia. How do you effectively govern a country that's home to one in five people on the planet, with an increasingly complex economy and society, if you don't allow public debate, civil activism, and electoral feedback? How do you gather enough information to actually make decisions? And how does a government that doesn't invite its citizens to participate still engender trust and bend public behavior without putting police on every doorstep?

Hu Jintao, China's leader from 2002 to 2012, had attempted to solve these problems by permitting a modest democratic thaw, allowing avenues for grievances to reach the ruling class.

His successor, Xi Jinping, has reversed that trend. Instead, his strategy for understanding and responding to what is going on in a nation of 1.4 billion relies on a combination of surveillance, AI, and big data to monitor people's lives and behavior in minute detail.

It helps that a tumultuous couple of years in the world's democracies have made the Chinese political elite feel increasingly justified in shutting out voters. Developments such as Donald Trump's election, Brexit, the rise of far-right parties across Europe, and Rodrigo Duterte's reign of terror in the Philippines underscore what many critics see as the problems inherent in democracy, especially populism, instability, and precariously personalized leadership.

Since becoming general secretary of the Chinese Communist Party in 2012, Xi has laid out a raft of ambitious plans for the country, many of them rooted in technology—including a goal to become the world leader in artificial intelligence by 2030. Xi has called for "cyber sovereignty" to enhance censorship and assert full control over the domestic internet. In May, he told a meeting of the Chinese Academy of Sciences that technology was the key to achieving "the great goal of building a socialist and modernized nation." In January, when he addressed the nation on television, the bookshelves on either side of him contained both classic titles such as *Das Kapital* and a few new additions, including two books about artificial intelligence: Pedro Domingos's *The Master Algorithm* and Brett King's *Augmented: Life in the Smart Lane.*

"No government has a more ambitious and far-reaching plan to harness the power of data to change the way it governs than the Chinese government," says Martin Chorzempa of the Peterson Institute for International Economics in Washington, DC. Even some foreign observers, watching from afar, may be tempted to wonder if such data-driven governance offers a viable alternative to the increasingly dysfunctional-looking electoral model. But overrelying on the wisdom of technology and data carries its own risks.

Data instead of dialogue

Chinese leaders have long wanted to tap public sentiment without opening the door to heated debate and criticism of the authorities. For most of imperial and modern Chinese history, there has been a tradition of disgruntled people from the countryside traveling to Beijing and staging small demonstrations as public "petitioners." The thinking was that if local authorities didn't understand or care about their grievances, the emperor might show better judgment.

Under Hu Jintao, some members of the Communist Party saw a limited openness as a possible way to expose and fix certain kinds of problems. Blogs, anticorruption journalists, human-rights lawyers, and online critics spotlighting local corruption drove public debate toward the end of Hu's reign. Early in his term, Xi received a daily briefing of public concerns and disturbances scraped from social media, according to a former US official with knowledge of the matter. In recent years, petitioners have come to the capital to draw attention to scandals such as illegal land seizures by local authorities and contaminated milk powder.

But police are increasingly stopping petitioners from ever reaching Beijing. "Now trains require national IDs to purchase tickets, which makes it easy for the authorities to identify potential 'troublemakers' such as those who have protested against the government in the past," says Maya Wang, senior China researcher for Human Rights Watch. "Several petitioners told us they have been stopped at train platforms." The bloggers, activists, and lawyers are also being systematically silenced or imprisoned, as if data can give the government the same information without any of the fiddly problems of freedom.

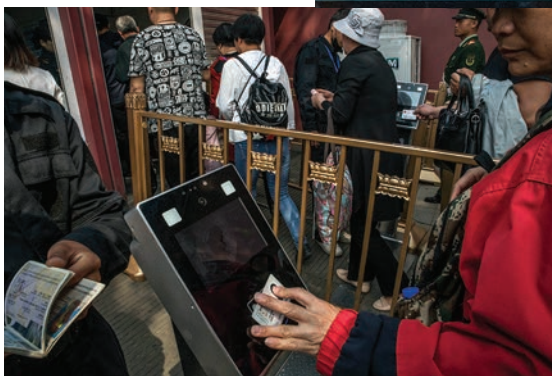

The idea of using networked technology as a tool of governance in China goes back to at least the mid-1980s. As Harvard political scientist Julian Gewirtz explains, "When the Chinese government saw that information technology was becoming a part of daily life, it realized it would have a powerful new tool for both gathering information and controlling culture, for making Chinese people more 'modern' and more 'governable'—which have been perennial obsessions of the leadership." Subsequent advances, including progress in AI and faster processors, have brought that vision closer.

As far as we know, there is no single master blueprint linking technology and governance in China. But there are several initiatives that share a common strategy of harvesting data about people and companies to inform decision-making and create systems of incentives and punishments to influence behavior. These initiatives include the State Council's 2014 "Social Credit System," the 2016 Cybersecurity Law, various local-level and private-enterprise experiments

Top: A Shanghai startup's demo of its system for facial recognition. Left: Visitors to Tiananmen Square in Beijing scan their IDs at a checkpoint.

in "social credit," "smart city" plans, and technology-driven policing in the western region of Xinjiang. Often they involve partnerships between the government and China's tech companies.

The most far-reaching is the Social Credit System, though a better translation in English might be the "trust" or "reputation" system. The government plan, which covers both people and businesses, lists among its goals the "construction of sincerity in government affairs, commercial sincerity, and judicial credibility." ("Everybody in China has an auntie who's been swindled. There is a legitimate need to address a breakdown in public trust," says Paul Triolo, head of the geotechnology practice at the consultancy Eurasia Group.) To date, it's a work in progress, though various pilots preview how it might work in 2020, when it is supposed to be fully implemented.

Blacklists are the system's first tool. For the past five years, China's court system has published the names of people who haven't paid fines or complied with judgments. Under new social-credit regulations, this list is shared with various businesses and government agencies. People on the list have found themselves blocked from borrowing money, booking flights, and staying at luxury hotels. China's national transport companies have created additional blacklists, to punish riders for behavior like blocking train doors or picking fights during a journey; offenders are barred from future ticket purchases for six or 12 months. Earlier this year, Beijing debuted a series of blacklists to prohibit "dishonest" enterprises from being awarded future government contracts or land grants.

A few local governments have experimented with social-credit "scores," though it's not clear if they will be part of the national plan. The northern city of Rongcheng, for example, assigns a score to each of its 740,000 residents, *Foreign Policy* reported. Everyone begins with 1,000 points. If you donate to a charity or win a government award, you gain points; if you violate a traffic law, such as by driving drunk or speeding through a crosswalk, you lose points. People with good scores can earn discounts on winter heating supplies or get better terms on mortgages; those with bad scores may lose access to bank loans or promotions in government jobs. City Hall showcases posters of local role models, who have exhibited "virtue" and earned high scores.

"The idea of social credit is to monitor and manage how people and institutions behave," says Samantha Hoffman of the Mercator Institute for China Studies in Berlin. "Once a violation is recorded in one part of the system, it can trigger responses in other parts of the system. It's a concept designed to support both economic development and social management, and it's inherently political." Some parallels to parts of China's blueprint already exist in the US: a bad credit score can prevent you from taking out a home loan, while a felony conviction suspends or annuls your right to vote, for example. "But they're not all connected in the same way—there's no overarching plan," Hoffman points out.

One of the biggest concerns is that because China lacks an independent judiciary, citizens have no recourse for disputing false or inaccurate allegations. Some have found their names added to travel blacklists without notification after a court decision. Petitioners and investigative journalists are monitored according to another system, and people who've entered drug rehab are watched by yet a different monitoring system.

"Theoretically the drug-user databases are supposed to erase names after five or seven years, but I've seen lots of cases where that didn't happen," says Wang of Human Rights Watch. "It's immensely difficult to ever take yourself off any of these lists."

Occasional bursts of rage online point to public resentment. News that a student had been turned down by a college because of her father's inclusion on a credit blacklist recently lit a wildfire of online anger. The college's decision hadn't been officially sanctioned or ordered by the government. Rather, in their enthusiasm to support the new policies, school administrators had simply taken them to what they saw as the logical conclusion.

The opacity of the system makes it difficult to evaluate how effective experiments like Rongcheng's are. The party has squeezed out almost all critical voices since 2012, and the risks of challenging the system—even in relatively small ways—have grown. What information is available is deeply flawed; systematic falsification of data on everything from GDP growth to hydropower use pervades Chinese government statistics. Australian National University researcher Borge Bakken estimates that official crime figures, which the government has a clear incentive to downplay, may represent as little as 2.5 percent of all criminal behavior.

In theory, data-driven governance could help fix these issues—circumventing distortions to allow the central government to gather information directly. That's been the idea behind, for instance, introducing air-quality monitors that send data back to central authorities rather than relying on local officials who may be in the pocket of polluting industries. But many aspects of good governance are too complicated to allow that kind of direct monitoring and instead rely on data entered by those same local officials.

However, the Chinese government rarely releases performance data that outsiders might use to evaluate these systems. Take the cameras that are used to identify and shame jaywalkers in some cities by projecting their faces on public billboards, as well as to track the prayer habits of Muslims in western China. Their accuracy remains in question: in particular, how well can facial-recognition software trained on Han Chinese faces recognize members of Eurasian minority groups? Moreover, even if the data collection is accurate, how will the government use such information to direct or thwart future behavior? Police algorithms that predict who is likely to become a criminal are not open to public scrutiny, nor are statistics that would show whether crime or terrorism has grown or diminished. (For example, in the western region of Xinjiang, the available information shows only that the number of people taken into police custody has shot up dramatically, rising 731 percent from 2016 to 2017.)

"It's not the technology that created the policies, but technology greatly expands the kinds of data that the Chinese government can collect on individuals," says Richard McGregor, a senior fellow at the Lowy Institute and the author of *The Party: The Secret World of China's Communist Rulers*. "The internet in China acts as a real-time, privately run digital intelligence service."

In the city of Xiangyang, cameras linked to face-recognition technology project photos of jaywalkers, with names and ID numbers, on a billboard.

Algorithmic policing

Writing in the *Washington Post* earlier this year, Xiao Qiang, a professor of communications at the

University of California, Berkeley, dubbed China's data-enhanced governance "a digital totalitarian state." The dystopian aspects are most obviously on display in western China.

Xinjiang ("New Territory") is the traditional home of a Chinese Muslim minority known as Uighurs. As large numbers of Han Chinese migrants have settled in—some say "colonized"—the region, the work and religious opportunities afforded to the local Uighur population have diminished. One result has been an uptick in violence in which both Han and Uighur have been targeted, including a 2009 riot in the capital city of Urumqi, when a reported 200 people died. The government's response to rising tensions has not been to hold public forums to solicit views or policy advice. Instead, the state is using data collection and algorithms to determine who is "likely" to commit future acts of violence or defiance.

The Xinjiang government employed a private company to design the predictive algorithms that assess various data streams. There's no public record or accountability for how these calculations are built or weighted. "The people living under this system generally don't even know what the rules are," says Rian Thum, an anthropologist at Loyola University who studies Xinjiang and who has seen government procurement notices that were issued in building the system.

The algorithm is thought to highlight suspicious behaviors such as visiting a mosque or owning too many books.

In the western city of Kashgar, many of the family homes and shops on main streets are now boarded up, and the public squares are empty. When I visited in 2013, it was clear that Kashgar was already a segregated city—the Han and Uighur populations lived and worked in distinct sections of town. But in the evenings, it was also a lively and often noisy place, where the sounds of the call to prayer intermingled with dance music from local clubs and the conversations of old men sitting out late in plastic chairs on patios. Today the city is eerily quiet; neighborhood public life has virtually vanished. Emily Feng, a journalist for the *Financial Times*, visited Kashgar in June and posted photos on Twitter of the newly vacant streets.

The reason is that by some estimates more than one in 10 Uighur and Kazakh adults in Xinjiang have been sent to barbed-wire-ringed "reeducation camps"—and those who remain at large are fearful.

In the last two years thousands of checkpoints have been set up at which passersby must present both their face and their national ID card to proceed on a highway, enter a mosque, or visit a shopping mall. Uighurs are required to install government-designed tracking apps on their smartphones, which monitor their online contacts and the web pages they've visited. Police officers visit local homes regularly to collect further data on things like how many people live in the household, what their relationships with their neighbors are like, how many times people pray daily, whether they have traveled abroad, and what books they have.

All these data streams are fed into Xinjiang's public security system, along with other records capturing information on everything from banking history to family planning. "The computer program aggregates all the data from these different sources and flags those who might become 'a threat' to authorities," says Wang. Though the precise algorithm is unknown, it's believed that it may highlight behaviors such as visiting a particular mosque, owning a lot of books, buying a large quantity of gasoline, or receiving phone calls or email from contacts abroad. People it flags are visited by police, who may take them into custody and put them in prison or in reeducation camps without any formal charges.

Adrian Zenz, a political scientist at the European School of Culture and Theology in Korntal, Germany, calculates that the internment rate for minorities in Xinjiang may be as high as 11.5 percent of the adult population. These camps are designed to instill patriotism and make people unlearn religious beliefs. (New procurement notices for cremation security guards seem to indicate that the government is also trying to stamp out traditional Muslim burial practices in the region.)

While Xinjiang represents one draconian extreme, elsewhere in China citizens are beginning to push back against some kinds of surveillance. An internet company that streamed closed-circuit TV footage online shut down those broadcasts after a public outcry. The city of Shanghai recently issued regulations to allow people to dispute incorrect information used to compile social-credit records. "There are rising demands for privacy from Chinese internet users," says Samm Sacks, a senior fellow in the Technology Policy Program at CSIS in New York. "It's not quite the free-for-all that it's made out to be."

Christina Larson is an award-winning foreign correspondent and science journalist, writing mostly about China and Asia.

The data lobbying

By Andrew Zaleski

Photographs by Jared Soares

FiscalNote takes the intuition out of politics. Does it take the democracy out, too?

lord
of

Sue Zoldak is a public relations expert with a fierce competitive streak. Her surname in Slovak, as she likes to point out, means "mercenary." Her firm, the Zoldak Agency, uses targeted advertising and grassroots campaigning to help clients spur voters to press elected officials into voting yes or no on specific bills. While not strictly a lobbyist—she doesn't communicate directly with lawmakers—Zoldak fits squarely into the influence-peddling milieu of Washington, DC, with 15 years' experience on K Street, where lobbying firms are traditionally headquartered. Put simply, she's a go-to person for companies and organizations determined to shape public policy.

Lately Zoldak's been getting help from a new source—a data intelligence platform called FiscalNote, founded by a 26-year-old political whiz, Tim Hwang. For a current client in the health-care industry—which Zoldak declines to name—she's tracking states that want to amend their "certificate of need" laws. These obscure laws, which were mandated by Congress in 1974, require health-care companies to prove to state regulators that a community needs their new hospital, nursing home, or rehab clinic. The initial idea was that a local market could support only so many health facilities. If there were too many, and one ended up with empty beds, it would raise prices to cover its fixed costs, overcharging patients.

Unsurprisingly, the laws have occasioned a political power struggle. Lobbyists for upstart clinics and hospitals wine and dine state regulators to overturn them, while those working for established hospital groups lobby to keep them in place. The political jockeying has been so intense that Congress repealed the federal mandate in 1987. Since then, 14 states have jettisoned the laws, and more may follow suit.

Zoldak's client, which she describes as a coalition of think tanks and "impacted parties," is seeking to get certificate-of-need laws off the books and wants to know which statehouses are considering dumping or amending them. It's a massive job. Zoldak needs to know how state lawmakers have voted on such laws in the past, which companies have tried to influence them, how successful that influence has been, and how every final vote has gone. She can then take the data to her client.

Clever as Zoldak might be, her agency is a boutique firm. She doesn't have an army of staff to work the phones and comb through state records. Enter FiscalNote. One click and the platform shows the text of bills, along with their sponsors and cosponsors. Another click and it summarizes everything there is to know about the state legislators who could prop up or nix the rules: their voting histories, the frequency with which bills they sponsor become law, their effectiveness by topic (health, education, housing), their ideological views on different issues. After crunching the data, FiscalNote can predict how each one will vote. "That tells us whether or not we should be targeting specific districts with our message," Zoldak says.

Zoldak has been so impressed with FiscalNote that she has invited Hwang to speak at George Washington University's Graduate School of Political Management, where she's an adjunct professor. "A lot of people say they're going to disrupt lobbying. Tim is one of those people who actually has the potential to do that in the long run," she says. "He's the closest that we have to a Mark Zuckerberg walking around."

Actually, the best analogy for FiscalNote may be not Facebook but *Moneyball*. Lobbying, like baseball, no longer belongs to old-timers and their seasoned intuition: it is now being refined by computer data and forecasts. There are other new digital players in town, including PopVox and Quorum. But FiscalNote—which has 1,300 clients and is backed by $50 million in venture funding from top investors like Mark Cuban, Steve Case, and Jerry Yang—is the marquee name.

"There is always going to be a personal touch in this business," says John Runyan, a longtime DC lobbyist for corporations and the president of Runyan Public Affairs, an independent government relations firm. But, he says, lobbyists using techniques made by a platform like FiscalNote "can really bore in

on exactly where they need to be focusing" to sway a legislator.

FiscalNote likes to say it represents a new force for democracy, putting the power of government data and analysis in the hands of the little guys: teachers' unions, environmental groups, and nonprofits of many stripes. But DC analysts contend that it's simultaneously helping entrenched lobbyists for candidates and multinational corporations to refine their targeted messages, potentially undercutting the will of the voters.

Weeding out the fakes

Hwang wears dark, rectangular glasses that echo his square face and possesses a self-assured demeanor that belies his age—though he does tend to punctuate sentences with a rhetorical-sounding "Right?" His ambitions, however, are nothing if not grown-up. "Our goal is to try and create a technology platform that aggregates every law and every regulation that governs all of humanity in every country on the planet," he told me this spring.

Hwang's fascination with technology in politics began in high school. He grew up in Montgomery County, Maryland, in a high-class suburb just across the line from DC, and volunteered for Barack Obama's 2008 presidential campaign when he was 15 years old. One of his jobs was organizing busloads of canvassers from reliably Democratic Maryland to visit key battleground precincts in Virginia.

"The culture of the Obama campaign was very startup-oriented, very decentralized, and very metrics-driven," he says. "We were thinking about how to leverage field data on a day-to-day basis in terms of aligning resources."

After high school Hwang attended Princeton University, where he took courses in algorithms, statistics, and public policy. In his senior year, he built an automated script that aggregated the privacy policies of the top 1,000 websites as ranked by Amazon's web analytics company, Alexa (not to be confused with its digital assistant). He found that most websites weren't complying with privacy regulations. That discovery became the germ for FiscalNote, and in 2013, at age 21, he went west with three colleagues to try to bootstrap a company. Too poor to pay San Francisco rents, the four ended up sharing a $70-a-night Motel 6 room. "We honestly didn't know what the hell we were doing," he says. Today their office commands the sixth floor of an office building on Pennsylvania Avenue, between the White House and the US Capitol.

With FiscalNote, Hwang wanted to make government data more useful by organizing it so that companies and organizations—and their lobbyists—could better predict what new laws and regulations would mean for them. As he told *Politico* in January, the company's technology can "enable the top attorney at McDonald's to immediately understand every single law and regulation pertaining to their industry."

The cleverness of his platform lies in the way it synthesizes myriad sources of data: hundreds of government websites, the text of reports published by the Congressional Budget Office and the Congressional Research Service, the rulings of regulatory agencies, lawmakers' voting records, and much more. A 46-person research and development team—split among squads tasked with data ingestion, data management, web applications, development operations, quality assurance, and product management—is constantly scraping the web, grabbing data sets, organizing them, and structuring them to make the text searchable. The scale is daunting: contact information for more than 78,000 elected officials (and their staffs) worldwide, public policy documents from 22 countries, and every regulation from every US regulatory agency going back 110 years.

But simply gaining a consolidated source of information isn't why groups from Toyota to the American Society for the Prevention of Cruelty to Animals pay anywhere from $10,000 to several hundred thousand dollars a year for access. FiscalNote can gauge a bill's chances of becoming law. Its algorithms can also assess how effective individual legislators are, as measured by legislative accomplishments, the party makeup of the legislature, and whether the bills they've backed were aimed at changing laws or were merely resolutions, memorials, or commendations. "Anyone can get data. It's really much more about

"He's the closest that we have to a Mark Zuckerberg walking around."

Hwang volunteered for Obama's 2008 campaign at age 15 and started FiscalNote at 21.

connecting all of the disparate data that has traditionally existed and helping you derive insights from it," says Gerald Kierce-Iturrioz, the company's product marketing manager.

That kind of insight can make lobbyists better at their jobs, says Rebecca Mark, a former congressional staffer who was most recently head of public affairs and policy at Cruise Automation, a company developing autonomous vehicles. To convince policymakers to support a proposal, "you have to explain why it's good for business, why it's good for the American public, and why it's good for the policymaker," she says. But lobbyists don't necessarily have access to hard data to back their claims. "That's why a tool like FiscalNote will make it easier and more efficient to do that job," Mark says.

FiscalNote showed off its data skills in 2017, when it analyzed each of the 22 million comments made on the Federal Communications Commission's website about the agency's proposed plan to repeal net neutrality. The company determined that 19 million commenters opposed the repeal. But it also discovered that hundreds of thousands of pro-repeal comments were written by bots using natural-language generation, an artificial-intelligence technique that simulates human language. Using its own tools for natural-language analysis, FiscalNote showed that each fraudulent comment consisted of 35 phrases arranged in the same order but varied by plugging in up to 25 interchangeable words and phrases, a system designed to make comments appear unique. It wasn't just net neutrality's opponents who relied on automation, though; FiscalNote also discovered thousands of pro-neutrality comments that turned out to be auto-generated letters,

with slight variations, selected by visitors to a website created by the Electronic Frontier Foundation, a digital-rights advocacy group. In a blog post, the company said the debate over net neutrality "serves as a prominent warning that, soon enough, the distinction between human- and computer-generated language may be nearly impossible to draw." FiscalNote, Hwang says, was able to make that distinction almost instantaneously.

FiscalNote can also alert clients to laws or proposals that will affect them, Hwang says. This spring, analysts at Southwest Airlines, using the platform, learned of a legislative meeting in which the airport authority in the state of Rhode Island would argue for a bill to tax fuel bought at T.F. Green, the state's largest airport. Southwest dispatched representatives to the state, where airlines and an airline advocacy association managed to beat back the bill.

Hwang says that a future version of FiscalNote can assimilate what clients have been doing on the platform and even recommend new political strategies. "I don't think we've tapped into the true potential of the work we do just yet," he says.

Aiding obstruction?

So does a service like FiscalNote make lobbying more egalitarian or less? It depends whom you ask.

Hwang portrays it as a way to level the playing field in politics. "Now, whether you're a local union or a mega-corporation, you can get the same information around the ideology of lawmakers," he says. "It's a different world in which the tools and the information that used to only be available for the wealthiest, most connected lobbyists and politicians and organizations can now be distributed across the world."

This view gets support from Christian Hoehner, the policy director of the Data Coalition, a K Street trade association that lobbies to make government data more readily available and transparent. Hoehner is a fan of FiscalNote, and uses it in his job. "It helps us figure out who the key players are, start tracking bills, and set up alerts," Hoehner says. "If you didn't use FiscalNote, you'd lose the ability to quickly find members of Congress and their staff representatives. At a very high level, FiscalNote helps democratize the government affairs function. It helps a small team be effective."

Lorelei Kelly, however, is not quite as sanguine. She runs the Resilient Democracy Coalition at Georgetown University's Beeck Center, where she studies how Congress can function better in the digital era. Even for small groups, she points out, FiscalNote costs several thousand dollars. "It optimizes certain information for people who can afford it," she says. "But the cost for participating in democracy should be zero. So unless something like this is available to citizens, it's not democratic."

Tim LaPira, a political science professor at James Madison University, takes a similar view. A former researcher for the Center for Responsive Politics, where he created a database on lobbying now available on OpenSecrets.org, he argues that FiscalNote will help the strong get stronger: it puts a wealth of information in powerful hands, making it easier for lobbyists to zero in on targets and protect their clients' preferred positions. The real work of lobbying, LaPira says, involves tireless efforts not to usher in new laws but to prevent old ones from changing. FiscalNote, he adds, is going to aid that kind of obstruction "more than it's going to help the little guy get something to happen."

The criticisms rankle Hwang. The business of Washington is always and inherently human, he says, and FiscalNote can't replace that human work; it only provides data. How that data is used is up to his clients. "To be in Washington is to have an opinion, and a tool like FiscalNote is a pretty substantial weapon to advance your agenda," he says. As more data goes digital, he adds, the machinations of how policy is made will become more transparent. In Hwang's version of Washington, even the smallest voices will be able to be heard over the din of the political class. And when that happens, the desires of voters, not the forces of technology, will shape political outcomes. Right?

Andrew Zaleski is a writer based near Washington, DC, who covers science, technology, and business.

"It optimizes certain information for people who can afford it. But the cost for participating in democracy should be zero."

WORLD-CLASS EXECUTIVE EDUCATION AT BROWN UNIVERSITY

Transforming Mid-Career Professionals

An Executive Master's degree from Brown University will prepare you to lead your organization, transform your field, and build a powerful lifelong professional network.

You will join a vibrant learning community and apply your new knowledge and skills through a critical challenge project.

The results can be seen in our alumni who are proven leaders; impacting their organizations and the world.

BROWN
School of Professional Studies

brown.edu/professional

IE Brown Executive MBA

Executive Master in Cybersecurity

Executive Master in Science and Technology Leadership

Executive Master of Healthcare Leadership

What comes next

Some consultants would just like a quick look inside your brain (p64). Alexa, please teach me how to spread propaganda (p72). Taiwan is testing out the future of democracy (p74). A science fiction vision of where the fight against fake news will ultimately take us (p80). And if you're not sure whether you're talking to a bot, here's a handy flowchart (p88).

3

NEUROPOLITICS

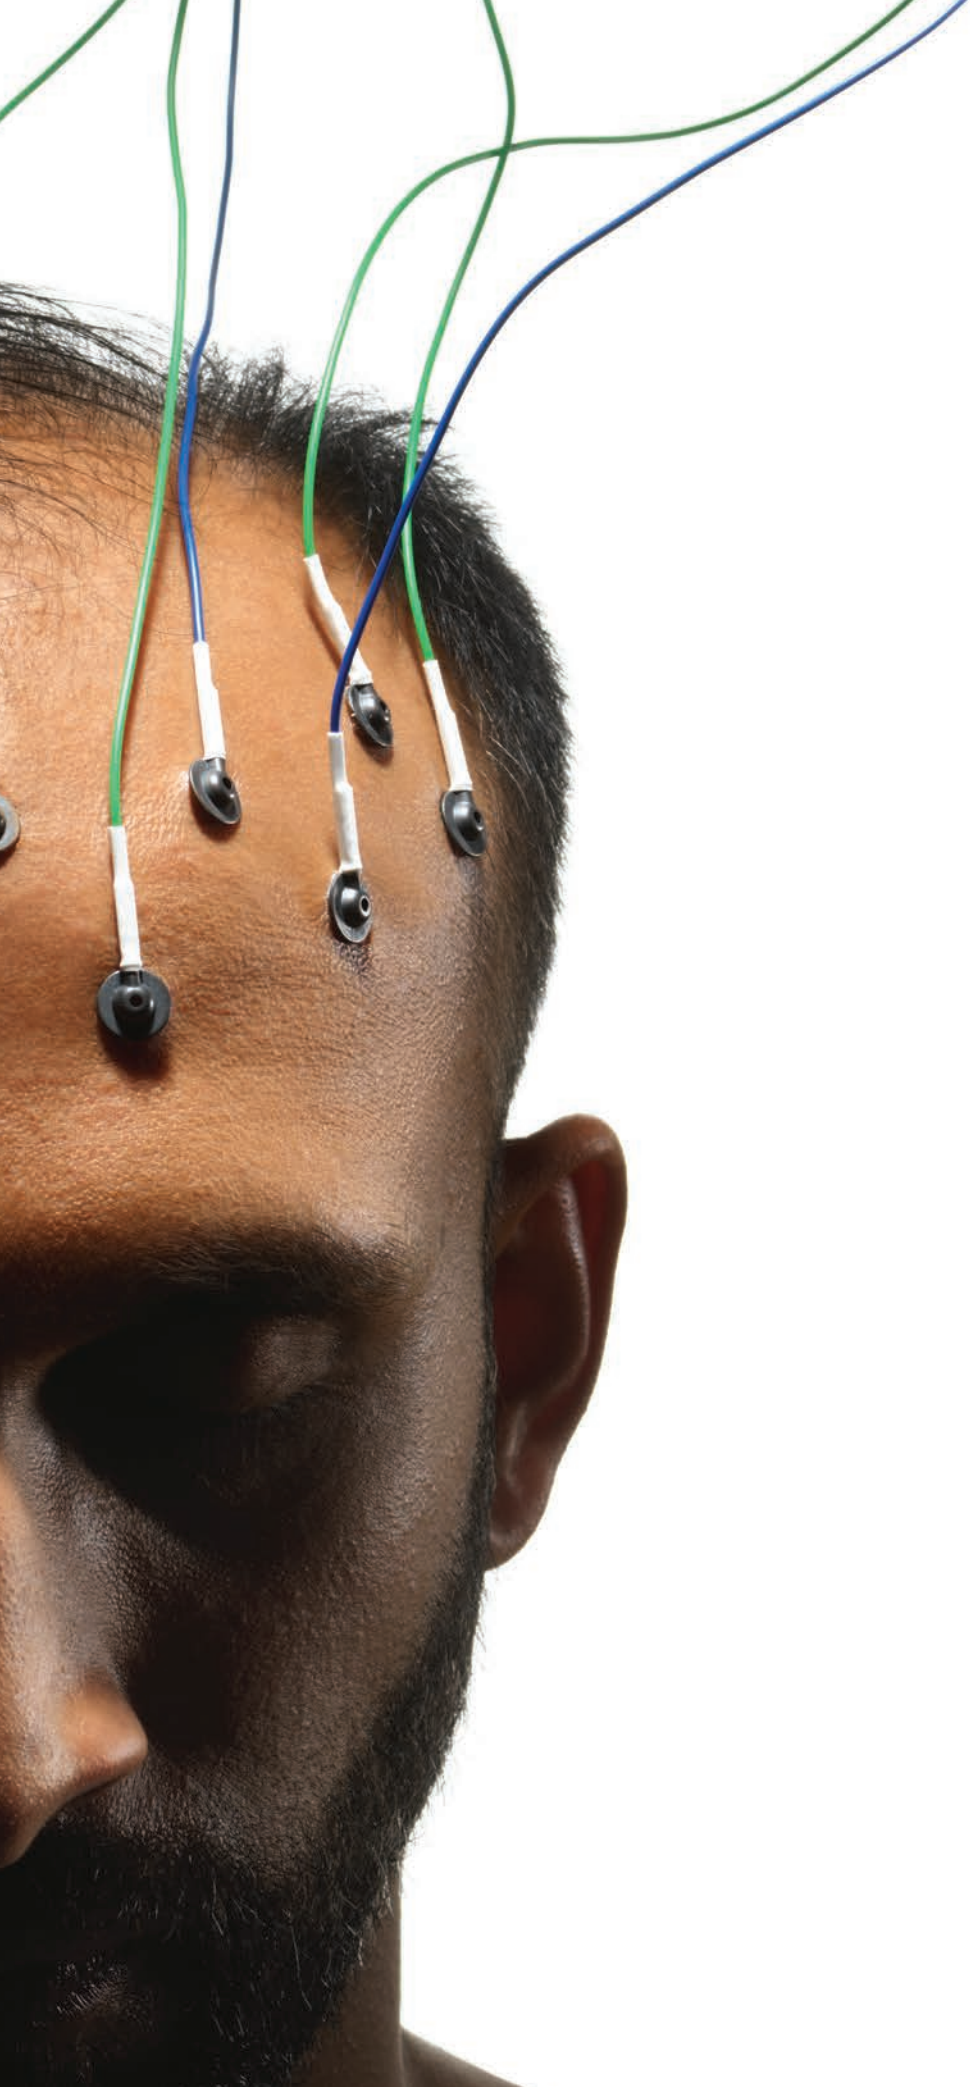

Meet the consultants who divine your political preferences by peering inside your brain.

BY ELIZABETH SVOBODA

PHOTO BY BRUCE PETERSON

Maria Pocovi slides her laptop over to me with the webcam switched on. My face stares back at me, overlaid with a grid of white lines that map the contours of my expression. Next to it is a shaded window that tracks six "core emotions": happiness, surprise, disgust, fear, anger, and sadness. Each time my expression shifts, a measurement bar next to each emotion fluctuates, as if my feelings were an audio signal. After a few seconds, a bold green word flashes in the window: ANXIETY. When I look back at Pocovi, I get the sense she knows exactly what I'm thinking with one glance.

Petite with a welcoming smile, Pocovi, the founder of Emotion Research Lab in Valencia, Spain, is a global entrepreneur par excellence. When she comes to Silicon Valley, she doesn't even rent an office—she just grabs a table here at the Plug and Play coworking space in Sunnyvale, California. But the technology she's showing me is at the forefront of a quiet political revolution. Campaigns around the world are employing Emotion Research Lab and other marketers versed in neuroscience to penetrate voters' unspoken feelings.

This spring there was a widespread outcry when American Facebook users found out that information they had posted on the social network—including their likes, interests, and political preferences—had been mined by the voter-targeting firm Cambridge Analytica. While it's not clear how effective they were, the company's algorithms may have helped fuel Donald Trump's come-from-behind victory in 2016.

But to ambitious data scientists like Pocovi, who has worked with major political parties in Latin America in recent elections, Cambridge Analytica, which shut down in May, was behind the curve. Where it gauged people's receptiveness to campaign messages by analyzing data they typed into Facebook, today's "neuropolitical" consultants say they can peg voters' feelings by observing their spontaneous responses: an electrical impulse from a key brain region, a split-second grimace, or a moment's hesitation as they ponder a question. The experts aim to divine voters' intent from signals they're not aware they're producing. A candidate's advisors can then attempt to use that biological data to influence voting decisions.

Political insiders say campaigns are buying into this prospect in increasing numbers, even if they're reluctant to acknowledge it. "It's rare that a campaign would admit to using neuromarketing techniques—though it's quite likely the well-funded campaigns are," says Roger Dooley, a consultant and author of *Brainfluence: 100 Ways to Persuade and Convince Consumers with Neuromarketing*. While it's not certain the Trump or Clinton campaigns used neuromarketing in 2016, SCL—the parent firm of Cambridge Analytica, which worked for Trump—has reportedly used facial analysis to assess whether what voters said they felt about candidates was genuine.

But even if US campaigns won't admit to using neuromarketing, "they should be interested in it, because politics is a blood sport," says Dan Hill, an American expert in facial-expression coding who advised Mexican president Enrique Peña Nieto's 2012 election campaign. Fred Davis, a Republican strategist whose clients have included George W. Bush, John McCain, and Elizabeth Dole, says that while uptake of these technologies is somewhat limited in the US, campaigns would use neuromarketing if they thought it would give them an edge. "There's nothing more important to a politician than winning," he says.

The trend raises a torrent of questions in the run-up to the 2018 midterms. How well can consultants like these use neurological data to target or sway voters? And if they are as good at it as they claim, can we trust that our political decisions are truly our own? Will democracy itself start to feel the squeeze?

Unspoken truths

Brain, eye, and face scans that tease out people's true desires might seem dystopian. But they're offshoots of a long-standing political tradition: hitting voters right in the feels. For more than a decade, campaigns have been scanning databases of consumer preferences—what music people listen to, what magazines they read—and, with the help of computer algorithms, using that information to target appeals to them. If an algorithm shows that middle-aged female SUV drivers are likely to vote Republican and care about education, chances are they'll receive campaign messages crafted explicitly to push those buttons.

Biometric practitioners say they can tap into truths that voters are unable to express.

Biometric technologies raise the stakes further. Practitioners say they can tap into truths that voters are often unwilling or unable to express. Neuroconsultants love to cite psychologist Daniel Kahneman, winner of the Nobel Prize in economics, who distinguishes between "System 1" and "System 2" thinking. System 1 "operates automatically and quickly, with little or no effort and no sense of voluntary control," he writes; System 2 involves conscious deliberation and takes longer.

"Before, everyone was focused on System 2," explains Rafal Ohme, a Polish psychologist who says his firm, Neurohm, has advised political campaigns in Europe and the United States. For the past decade, Ohme has devoted most of his efforts to probing consumers' and voters' System 1 leanings, which he thinks is as important as listening to what they say. It's been great for his business, he says, because his clients are impressed enough with the results to keep coming back for more.

Many neuroconsulting pioneers built their strategy around so-called "neurofocus groups." In these studies, involving anywhere from a dozen to a hundred people, technicians fit people's scalps with EEG electrodes and then show them video footage of a candidate or campaign ad. As subjects watch, scalp sensors pick up electrical impulses that reveal, second by second, which areas of the brain are activated.

ELIZABETH SVOBODA; PREVIOUS SPREAD: MODEL: SHAHJEHAN KAHN/MODEL CLUB INC; STYLIST: BRE WELCH

"One of the things we can analyze is the attentional process," says Mexico City neurophysiologist Jaime Romano Micha, whose former firm, Neuropolitka, was one of the top providers of brain-based services to political campaigns. Romano Micha would place electrodes on a subject's scalp to detect activity in the reticular formation, a part of the brain stem that tracks how engaged someone is. So if subjects are watching a political ad and activity in their reticular formation spikes, say, 15 seconds in, it means the message has truly caught their attention at that point.

Other brain areas provide important clues too, Romano Micha says. Electrical activity on the left side of the cerebral cortex suggests people are working hard to understand a political message; similar activity on the right side may reveal the precise moment the message's meaning clicks into place. With these kinds of insights, campaigns can refine a message to maximize its oomph: placing the most gripping moment at the beginning, for instance, or cutting the parts that cause people's attention to wander.

But while brain imaging remains part of the neuropolitical universe, most neuroconsultants say it's hardly sufficient by itself. "EEG gives us very general information about the decision process," Romano Micha says. "Some people are saying that through EEG we can go into the mind of people, and I think that's not possible yet." There are cheaper and more reliable tools, several consultants claim, for getting at a voter's true feelings and desires.

"There's nothing more important to a politician than winning."

Electrodes everywhere

EEG scans, in fact, are now just one in a smorgasbord of biometric techniques. Romano Micha also uses near-infrared eye trackers and electrodes around the orbital bone to track "saccades," minuscule movements of the eye that indicate viewers' attentional focus as they watch a campaign spot. Other electrodes supply a rough gauge of arousal by measuring electrical activity on the surface of a person's skin.

Of course, you can't stick electrodes on every person watching TV and browsing Facebook. But you don't need to. The results from experiments on small neuro-focus groups can be used to influence voters who aren't being sampled themselves. If, for example, biodata reveals that liberal women over 50 are fearful when they see an ad about illegal immigration, campaigns that want to stoke such fear can broadcast that same message to millions of people with similar demographic and social profiles.

Pocovi's approach at Emotion Research Lab requires only a video player and a front-facing webcam. When volunteers enroll in her political focus groups online, she sends them videos of an ad spot or a candidate that they can watch on their laptop or phone. As they digest the content, she tracks their eye movements and subtle shifts in their facial expressions.

"We have developed algorithms to read the microexpressions in the face and translate in real time the emotions people are feeling," Pocovi says. "Many times, people tell you, 'I'm worried about the economy.' But what are really the things that move you? In my experience, it's not the biggest things. It's the small things that are close to you." Something as small as a candidate's inappropriately furrowed brow, she says, can color our perception without our realizing it.

Pocovi says her facial analysis software can detect and measure "six universal emotions, 101 secondary emotions, and eight moods," all of which interest campaigns anxious to learn how people are responding to a message or a candidate. She also offers a crowd-analytics service to track the emotional reactions of individual faces in a human sea, meaning that campaigns can take the temperature of a room as their candidate is speaking.

ERL's software is built around the facial action coding system (FACS) developed

The experts hope to discover voters' feelings via signals they're not even aware they're producing.

NEUROMARKETING S.A. DE C.V. AND DR. JAIME ROMANO MICHA

"I measure hesitation," says Ohme. "I can change your mind only if you hesitate. If you are a firm believer, I cannot change anything."

by Paul Ekman, a famed American psychologist. Pocovi's algorithm deconstructs each facial image from the webcam into more than 50 "action units," movements of specific muscle groups. Distinct clusters of action units correspond to particular emotions: cheek and outer-lip muscles contracting at the same time reveal happiness, while lowered brows and raised upper eyelids betray anger. Pocovi trains her system to recognize each one by showing it many reference images from a large database of faces expressing that emotion.

Some critics of Ekman's system, such as neuroscientist Lisa Feldman Barrett, have argued that facial expressions don't necessarily correlate with emotional states. Still, a variety of studies have shown at least some correspondence. In a 2014 study at Ohio State University, cognitive scientists defined 21 "distinct emotions," based on the consistent ways most of us move our facial muscles.

Pocovi says her surveys also operate as an image-refining tool for candidates themselves. She analyzes video of candidates to pinpoint precise moments when their expressions make voters feel confused, disgusted, or angry. Politicians can then use this information to rehearse a different emotional approach, which can itself be vetted using Pocovi's survey platform until it produces the desired response in viewers. In one campaign Pocovi advised, a candidate was recording a TV ad spot with an uplifting, positive message, but it kept getting terrible reviews in test screenings. The spot's poor performance was a mystery—until Pocovi's analysis of the candidate's face showed he was unwittingly conveying anger and disgust. Once he realized what was going on, he was able to tweak his presentation and get a better response from the public.

Several onetime devotees of brain-scan analysis are also pursuing simpler and cheaper techniques these days. Before the 2008 financial crisis, Ohme says, international clients were more willing to fly five guys from Poland out to perform on-site brain studies. After the recession, though, that business mostly dried up.

That prompted Ohme to develop a different strategy, one untethered to time, space, or EEG electrodes. His updated approach stems from that used in unconscious-bias studies by social psychologist Anthony Greenwald, who became

DMITRY KOSTYUKOV

a mentor when Ohme visited the US on a Fulbright scholarship. Ohme says his smartphone-based test—which he calls iCode—reveals covert political leanings that would never surface in traditional questionnaires or focus groups.

Ohme's survey takers begin by answering calibration questions to assess their baseline reaction time. A habitually slower person, for instance, might have a "unit time" lasting 585 milliseconds, while someone quicker might take 387 milliseconds. Then images of politicians are shown on the screen, each paired with a single attribute, such as "trustworthy," "well-known," or "shares my values." Users tap "yes" or "no" to indicate whether they agree with each pairing. As the test proceeds, the app tracks not just how they answer but how quickly they touch the screen and what tapping rhythm they establish.

What's interesting, Ohme says, isn't how people respond to the questions per se, but how much they dither first. "When we measure the hesitation level, we can see that some answers are positive but with hesitation, and some are positive and instantaneous," he says. "We measure how much you deviated [from baseline]. This deviation is key."

Ohme declines to discuss his current political clients in much detail, citing confidentiality agreements. But he volunteers that in an iCode survey of nearly 900 people, he predicted Hillary Clinton's 2016 defeat before the election. Throughout the year, Clinton ran comfortably ahead of Trump in traditional polls. But when Ohme asked test subjects whether Clinton shared their values, they often hesitated for an unusually long time before responding that she did. Ohme knew a sense of shared values was a big factor motivating people to vote in 2016 (in previous elections "powerful" and "leader" were key), so the results of the test gave him serious doubts about a Clinton victory. He argues that if Clinton's campaign had run one of his studies before the election, she would have understood the depth of her vulnerability and could have made course corrections.

Ohme claims to have helped other candidates in similar straits. One of his tests revealed that while a certain European client had a good-sized base of supporters, many weren't motivated to get out and vote because they assumed their candidate would win. Armed with this knowledge, the campaign made a renewed push to get its loyal base to the polls. The client ended up winning in a squeaker.

When Ohme asked test subjects whether Hillary Clinton shared their values, they often hesitated for an unusually long time.

The biggest lies in life

Does measuring people's spontaneous reactions to a TV ad or a stump speech tell you how they will ultimately vote, however? "On the applied side, it's pretty unclear, the hype from the reality," says Darren Schreiber, a professor in political science at the University of Exeter and author of *Your Brain Is Built for Politics*. "It's easy to over-believe the ability of these tools." So far cognitive tests have had mixed results. Contrasting studies have shown that implicit attitudes both do and don't predict how people vote.

Still, Schreiber, who has conducted brain-scan tests of political attitudes, admits the technologies are worrisome. Democracy assumes the presence of rational actors, capable of digesting information from all quarters and coming to reasoned conclusions. If neuroconsultants are even half as good as they claim at probing people's innermost thoughts and shifting their voting intentions, it calls that assumption into question.

"We are susceptible in multiple ways, and not aware of our susceptibility," Schreiber says. "The fact that attitudes can be manipulated in ways we're not aware of has a lot of implications for political discourse." If campaigns are nudging voters toward their candidate without voters' knowledge, political discussions that were once exchanges of reasoned views will become knee-jerk skirmishes veering ever further from the democratic ideal. "I don't think it's time to run in panic," Schreiber says, "but I don't think we can be sanguine about it."

Ohme insists that voters can inoculate themselves against neuroconsultants' tactics if they're savvy enough. "I measure hesitation. I can change your mind only if you hesitate. If you are a firm believer, I cannot change anything," he says. "If you're scared to be manipulated, learn. The more you learn, the more firm and stable your attitudes are, and the more difficult it is for someone to convince you otherwise."

That's perfectly reasonable advice. But I wonder. After meeting Pocovi, I logged into Emotion Research Lab to let its software track my face while I watched a demo video. The video was of a laughing baby, and I felt the corners of my mouth quirking up. After, the computer asked me how I'd felt while watching. "Happy," I clicked. I'm a mom, right? I love babies. Yet when my emotion analysis arrived, it showed almost no trace of happiness on my face.

Thinking about the results, I realized the emotion software was right. I hadn't really been happy at all. I had taken the test late at night, and I had been exhausted. The computer had seen me in a way I wasn't used to seeing myself. I thought of something Dan Hill, the former advisor to the Mexican president's campaign, had told me. "The biggest lies in life," he'd said, "are the ones we tell ourselves."

Elizabeth Svoboda is a science writer in San Jose, California, and the author of What Makes a Hero?: The Surprising Science of Selflessness.

MIT Technology Review

AI and robots are wreaking economic havoc. We need more of them.

The Economy Issue

Vol. 121 No. 4
Jul/Aug 2018
$9.99 USD
$10.99 CAD

Basic income could work—if you do it Canada-style

A Canadian province is giving people money with no strings attached—revealing both the appeal and the limitations of the idea.

Sustainable Energy

To Feed the World, Improve Photosynthesis

By reworking the basic metabolism of crops, plant scientists hope to forestall devastating food shortages.

by Katherine Bourzac

Rewriting Life

Inside the effort to print lungs and breathe life into them with stem cells

Martine Rothblatt wants to end transplant shortages with 3-D-printed lungs.

2018

35 Innovators Under 35

Meet Jonas Cleveland, CEO of COSY, and the other young inventors, pioneers, and entrepreneurs shaping the future for all of us.

The AI advances that brought you Alexa will be weaponized for political manipulation.

By **Lisa-Maria Neudert**

Teaching propaganda how to talk

The battle against propaganda bots is an arm's race for our democracy. It's one we may be about to lose. Bots—simple computer scripts—were originally designed to automate repetitive tasks like organizing content or conducting network maintenance, thus sparing humans hours of tedium. Companies and media outlets also use bots to operate social-media accounts, to instantly alert users of breaking news or promote newly published material.

But they can also be used to operate large numbers of fake accounts, which makes them ideal for manipulating people. Our research at the Computational Propaganda Project studies the myriad ways in which political bots employing big data and automation have been used to spread disinformation and distort online discourse.

Bots have proved to be one of the best ways to broadcast extremist viewpoints on social media, but also to amplify such views from other, genuine accounts by liking, sharing, retweeting, hearting, and following, just as a human would. By doing so they're gaming the algorithms, and rewarding the posts they've interacted with by giving them more visibility.

This will seem tame compared with what's on the way.

Lisa-Maria Neudert is a doctoral candidate at the Oxford Internet Institute and a researcher with the Computational Propaganda Project.

Strength in numbers

In the wake of Russia's interference in the 2016 US election came a wave of discussion about how to shield politics from propaganda. Twitter has taken down suspicious accounts, including bots, in the tens of millions this year, while regulators have proposed bot bans and transparency measures, and called for better cooperation with internet platforms.

So it may appear as if we're gaining the upper hand. And that's partly true—the bots' tactics have lost their novelty and never had finesse. Their strength used to lie in numbers. Propagandists would mobilize armies of them to flood the internet with posts and replies in an attempt to overwhelm genuine democratic discourse. As we've created technical countermeasures that are better at detecting bot-like behavior, it's become easier to shut them down. People, too, have become more alert and effective at spotting them. The average bot does little to conceal its robotic character, and a quick look at its patterns of tweeting, or even its profile picture, can give it away.

The next generation of bots is rapidly evolving, however. Owing in large part to advances in natural-language processing—the same technology that makes possible voice-operated interfaces like Amazon's Alexa, Google Assistant, and Microsoft's Cortana—these bots will behave a lot more like real people.

Admittedly, these conversational interfaces are still bumpy, but they're getting better, and the benefits of being able to successfully decode human language are tremendous. Digital assistants are just one use of them—brands operate conversational chatbots for customer service, and publishers like CNN use them to distribute personalized media content.

Such chatbots openly declare themselves to be automated, but the propaganda bots won't. They'll present themselves as human users participating in online conversation in comment sections, group chats, and message boards.

Contrary to popular belief, this isn't happening yet. Most bots merely react to keywords that trigger a boilerplate response,

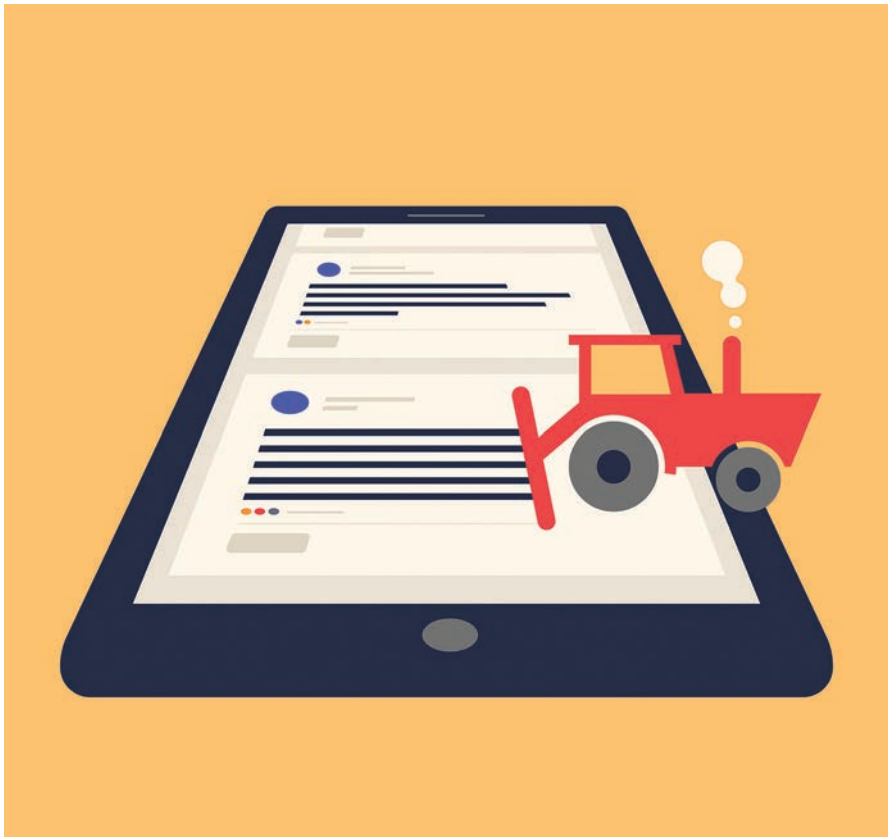

As bots learn how to understand context and intent, they become more adept at engaging in conversation without blowing their cover.

which rarely fits into the context or syntax of a given conversation. These responses are often easy to spot.

But it's getting harder. Already, some simple preprogrammed bot scripts have been successful at misleading users. As bots learn how to understand context and intent, they become more adept at engaging in conversation without blowing their cover.

In a few years, conversational bots might seek out susceptible users and approach them over private chat channels. They'll eloquently navigate conversations and analyze a user's data to deliver customized propaganda. Bots will point people toward extremist viewpoints and counter arguments in a conversational manner.

Rather than broadcasting propaganda to everyone, these bots will direct their activity at influential people or political dissidents. They'll attack individuals with scripted hate speech, overwhelm them with spam, or get their accounts shut down by reporting their content as abusive.

Great for Google, great for bots

It's worth taking a look at exactly how the AI techniques that power these kinds of bots are getting better, because the methods employed by tech companies also happen to be great for boosting the capabilities of political bots.

To work, natural-language processing requires substantial amounts of data. Tech companies like Google and Amazon get such data by opening their language-processing algorithms to the public via application programming interfaces, or APIs. Third parties—such as a bank, for example—that want to automate conversations with their customers can send raw data, such as the audio or text scripts of phone calls, to these APIs. Algorithms process the language and return machine-readable data ready to trigger commands. In return, the technology companies that provide these APIs get access to large amounts of conversational examples, which they can use to improve their machine learning and algorithms.

In addition, almost all major technology companies make open-source algorithms for natural-language processing available to developers. The developers can use these to build new, proprietary applications—software for a voice-controlled robot, for example. As developers advance and refine the original algorithms, the technology companies profit from their feedback.

The problem is that such services are widely accessible to almost anyone—including the people building political bots. By providing a toolkit for automating conversation, tech companies are unwittingly teaching propaganda to talk.

The worst is yet to come

Bots versed in human language remain outliers for now. It still requires substantial expertise, computing power, and training data to equip bots with state-of-the-art language-processing algorithms. But it's not out of reach. Since 2010 political parties and governments have spent more than half a billion dollars on social-media manipulation, turning it into a highly professionalized and well-funded sector.

There's still a long way to go before a bot will be able to spoof a human in one-on-one conversation. Yet as the algorithms evolve, those capabilities will emerge.

As with any other innovation, once these AI techniques are out of the box, they'll inevitably break free from the limited set of applications they were originally designed to perform.

vthe people

Can Taiwan's experiment in participatory lawmaking teach the world anything about the future of governing?

By CHRIS HORTON
Photographs by AN RONG XU


vTaiwan
歷史案件 5
遙控無人機應用與管理機制
Uncrewed Aircraft Vehicle / 無法管的無人機？
遙控無人機的管理與應用機制
共享經濟
Sharing Economy / 你的我的他的，什麼才是平
Airbnb自用住宅招租
Airbnb Short-Term Rental / 各式住宿類型，人

It was late in 2015, and things were at an impasse. Some four years earlier, Taiwan's finance ministry had decided to legalize online sales of alcohol. To help it shape the new rules, the ministry had kicked off talks with alcohol merchants, e-commerce platforms, and social groups worried that online sales would make it easy for children to buy liquor. But since then they had all been talking past each other. The regulation had gotten nowhere.

That was when a group of government officials and activists decided to take the question to a new online discussion platform called vTaiwan. Starting in early March 2016, about 450 citizens went to vtaiwan.tw, proposed solutions, and voted on them.

Within a matter of weeks, they had formulated a set of recommendations. Online alcohol sales would be limited to a handful of e-commerce platforms and distributors; transactions would be by credit card only; and purchases would be collected at convenience stores, making it nearly impossible for a child to surreptitiously get hold of booze. By late April the government had incorporated the suggestions into a draft bill that it sent to parliament.

The deadlock "resolved itself almost immediately," says Colin Megill, the CEO and cofounder of Pol.is, one of the digital platforms vTaiwan uses to host discussion. "The opposing sides had never had a chance to actually interact with each other's ideas. When they did, it became apparent that both sides were basically willing to give the opposing side what it wanted."

Three years after its founding, vTaiwan hasn't exactly taken Taiwanese politics by storm. It has been used to debate only a couple of dozen bills, and the government isn't required to heed the outcomes of those debates (though it may be if a new law passes later this year). But the system has proved useful in finding consensus on deadlocked issues such as the alcohol sales law, and its methods are now being applied to a larger consultation platform, called Join, that's being tried out in some local government settings. The question now is whether it can be used to settle bigger policy questions at a national level—and whether it could be a model for other countries.

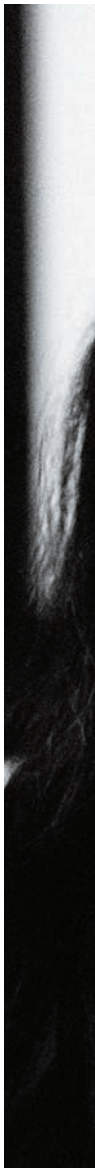

"The opposing sides had never had a chance to actually interact with each other's ideas."

Taiwan might not seem like the most obvious place for a pioneering exercise in digital democracy. The island held its first direct presidential election only in 1996, after a century marked first by Japanese colonial rule and then by Chinese nationalist martial law. But that oppressive past has also meant that Taiwanese have a history of taking to the streets to push back against heavy-handed government. In Taiwan's democratic era, it was a protest four years ago that planted the seed for this innovative political experiment.

The 2014 Sunflower Movement, led by students and activists, derailed an attempt by President Ma Ying-jeou's government to ram through a trade agreement between Taiwan, which has been locally ruled since 1949, and China, which claims Taiwan as its territory. For more than three weeks the protesters occupied government buildings over the deal, which they felt would give China too much leverage over the Taiwanese economy.

In the aftermath, the Ma government invited Sunflower activists to create a platform through which it might better communicate with Taiwan's youth. A Taiwanese civic tech community known as g0v (pronounced "Gov Zero"), which had played a leading role in the Sunflower protests, built vTaiwan in 2015 and still runs it. The platform enables citizens, civil-society organizations, experts, and elected representatives to discuss proposed laws via its website as well as in face-to-face meetings and hackathons. Its goal is to help policymakers make decisions that gain legitimacy through consultation.

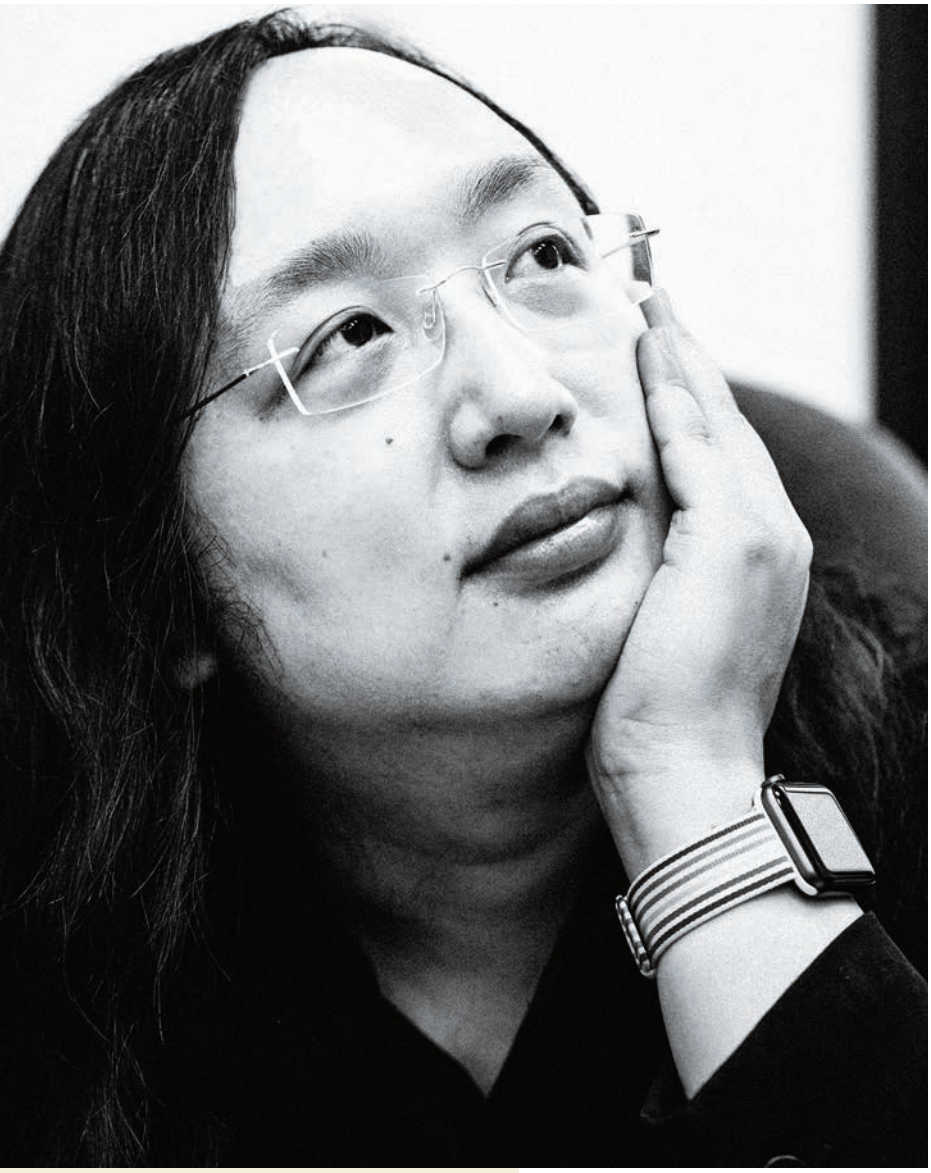

Wu Min Hsuan

Wu, an activist in the Sunflower Movement, says the government has missed out on chances to test vTaiwan on larger, non-digital issues such as pension reform. For vTaiwan to work, he says, "it needs real power."

Audrey Tang

Taiwan's digital minister was a leading hacker and Sunflower activist. She says that senior public servants need to see that people who comment online are "not protesters or mobs, but actually people with distinct expertise."

"I would say vTaiwan is about civil society learning the functions of the government and, to a degree, collaborating," Taiwan's digital minister, Audrey Tang, told me during a visit to her office. Tang, a famed hacker who helped the thousands of Sunflower protesters build and maintain their internal communications network, was appointed by the current president, Tsai Ing-wen, who won the 2016 election on a pledge of government transparency.

vTaiwan relies on a hodgepodge of open-source tools for soliciting proposals, sharing information, and holding polls, but one of the key parts is Pol.is, created by Megill and a couple of friends in Seattle after the events of Occupy Wall Street and the Arab Spring in 2011. On Pol.is, a topic is put up for debate. Anyone who creates an account can post comments on the topic, and can also upvote or downvote other people's comments.

That may sound much like any other online forum, but two things make Pol.is unusual. The first is that you cannot reply to comments. "If people can propose their ideas and comments but they cannot reply to each other, then it drastically reduces the motivation for trolls to troll," Tang says.

The second is that it uses the upvotes and downvotes to generate a kind of map of all the participants in the debate, clustering together people who have voted similarly. Although there may be hundreds or thousands of separate comments, like-minded groups rapidly emerge in this voting map, showing where there are divides and where there is consensus. People then naturally try to draft comments that will win votes from both sides of a divide, gradually eliminating the gaps.

"The visualization is very, very helpful," Tang says. "If you show people the face of the crowd, and if you take away the reply button, then people stop wasting time on the divisive statements."

In one of the platform's early successes, for example, the topic at issue was how to regulate the ride-hailing company Uber, which had—as in many places around the world—run into fierce opposition from local taxi drivers. As new people joined the online debate, they were shown and asked to vote on comments that ranged from calls to ban Uber or subject it to strict regulation, to calls to let the market decide, to more general statements such as "I think that Uber is a business model that can create flexible jobs."

Within a few days, the voting had coalesced to define two groups, one pro-Uber and one, about twice as large, anti-Uber. But then the magic happened: as the groups sought to attract more supporters, their members started posting comments on matters that everyone could agree were important, such as rider safety and liability insurance. Gradually, they refined them to garner more votes. The end result was a

PREVIOUS SPREAD: LAM YIK FEI (PROTEST)

set of seven comments that enjoyed almost universal approval, containing such recommendations as "The government should set up a fair regulatory regime," "Private passenger vehicles should be registered," and "It should be permissible for a for-hire driver to join multiple fleets and platforms." The divide between pro- and anti-Uber camps had been replaced by consensus on how to create a level playing field for Uber and the taxi firms, protect consumers, and create more competition. Tang herself took those suggestions into face-to-face talks with Uber, the taxi drivers, and experts, which led the government to adopt new regulations along the lines vTaiwan had produced.

vTaiwan's website boasts that as of August 2018, it had been used in 26 cases, with 80 percent resulting in "decisive government action." As well as inspiring regulations for Uber and for online alcohol sales, it has led to an act that creates a "fintech sandbox," a space for small-scale technological experiments within Taiwan's otherwise tightly regulated financial system.

"It's all solving the same problem: essentially saying, 'What if we're talking about things that are emergent, [for which] there are only a handful of early adopters?'" Tang says. "That's the basic problem we were solving at the very beginning with vTaiwan."

But while vTaiwan can bridge gulfs in public opinion, what it can't always overcome is politics. After the Tsai administration took office in 2016, it withdrew all bills awaiting legislative approval. Observers chalked that up to the new president's desire to differentiate her agenda from her predecessor's. The online alcohol sales bill that the Ma government had drafted from the vTaiwan suggestions never saw the light of day.

That the government isn't required to heed discussions on vTaiwan is the system's biggest shortcoming. Jason Hsu, a former activist and now opposition legislator who helped bring vTaiwan into being during the Ma administration, calls it "a tiger without teeth."

Moreover, the Tsai administration has chosen to use it only for issues, such as regulating Uber, that have to do with the digital economy. That's because people who care about such issues are the ones most likely to be comfortable using a digital discussion platform. But some think it won't get serious traction with the public unless it is put to use on non-digital issues that matter to more people. C.L. Kao, one of the cofounders of g0v, argues that the

Jason Hsu ↓

A former activist, and now an opposition legislator, Hsu helped bring the vTaiwan platform into being. He says its big flaw is that the goverment is not required to heed its discusions.

Karen Yu →

A lawmaker for the ruling party, Yu says that vTaiwan is "not a huge priority" for the government and has come "close to death" at times.

If vTaiwan's recommendations are ignored, the process runs the risk of being viewed as "openwashing."

government could have applied vTaiwan to two contentious recent issues, pension reform and labor reform, as a way to build its credibility.

In any case, Kao says, if vTaiwan's recommendations are ultimately ignored, as they were with the alcohol sales law, then the whole process runs the risk of being viewed as "openwashing"—something that creates the pretense of transparency. "The end goal is legislation," he says.

vTaiwan is one of dozens of participatory governance projects around the world listed on CrowdLaw, a site run by the Governance Lab at New York University. Most of them, says Beth Noveck, the lab's director, suffer from the same problem: they're not binding on governments, which means it's also hard for them to gain credibility with citizens. Still, she says, Taiwan's experiment is "a step in the right direction." It's "far more institutionalized" than what's been seen elsewhere, she adds.

The platform may be about to get a little more clout. This autumn legislators will debate and vote on a digital communications bill that, among other things, says that "digital-economy issues are to be deliberated in an open, multistakeholder process that the government has the duty to support," in Tang's words. But what "support" means—how much weight lawmakers or the government will have to give to vTaiwan's deliberations—is still up in the air.

Taiwan does have a newer participatory governance system that is getting more traction. Join, also overseen by Audrey Tang, is a platform for hosting and debating online petitions, again using Pol.is to create consensus. She describes it as a vTaiwan within the government—"basically the same … process, but with senior career public servants instead of g0v volunteers" at the heart of the platform.

Although petitions on Join still aren't legally binding, any government agency that agrees to participate in a deliberation must, if the petition gets more than 5,000 signatures, give a point-by-point response explaining why it agreed to or rejected the proposal. Five of Taiwan's cities or counties are testing Join; the aim is ultimately to roll it out nationwide, Tang says.

Join tends to attract a broader, older, and less tech-savvy range of users than vTaiwan. The advantage of this, says Tang, is that it doesn't tackle only digital-economy issues, as vTaiwan does, but a wide variety of questions, "like whether we should build a hospital in the southmost part of Taiwan, in Hengchun, or whether the first publicly open marine national park should ban fishing." The downside is that there's more resistance from the government bureaucracy. Senior public servants "need some hand-holding," she says, to be able to see people who comment online as "not protesters or mobs, but actually people with distinct expertise."

While only 200,000 people have so far taken part in a vTaiwan discussion, nearly five million of the country's 23 million inhabitants are already on Join. More than 10,000 voted on a recent proposal that advocated caning as a punishment for drunk driving, sexual assault, and child abuse.

Here too, the consensus-building tendencies of Pol.is can lead the discussion in unexpected directions. Initially, opinion on the caning issue was divided into three camps: besides the people who were for and against caning, a third group argued that it was too light a punishment for such offenses.

Eventually, however, the consensus opinions that emerged had nothing to do with caning at all, but were more focused on methods of preventing those crimes. At the time of this writing, proposals being considered for legislation included alcohol locks and confiscating drunk drivers' cars.

This suggests people had concluded that, in fact, "To cane or not to cane?" was the wrong question to ask. That kind of realization, and solution, wouldn't have emerged from a traditional online petition that only gives people the option of voting yes or no.

Karen Yu, a legislator in President Tsai's Democratic Progressive Party, says vTaiwan is "not a huge priority" for the administration and has come "close to death" at times. Join, she points out, at least benefits from the legitimacy of being managed by the government. Wu Min Hsuan, an activist who occupied Taiwan's Legislative Yuan during the Sunflower Movement protests, says Join has already proved itself much more productive than vTaiwan. The obstacle, he believes, is political will. "The experiment is important and has value," he says. "But the platform has its limits. It needs real power."

Chris Horton is a journalist based in Taipei. An Rong Xu is a photojournalist in New York City.

By Karl Schroeder

Illustrations by Rob Sheridan

Noon in the antilibrary

Marius cursed and jammed a mic stand between the crash bars of the TV studio door. "If SWAT's on its way, we don't have much time," he said.

"I don't understand." Michaela, who up until a couple of minutes ago had been streaming their interview live, still sat on one of the oval chairs under the hot lights. "What are they talking about?"

The cube-shaped television studio had black-painted walls surrounding the bright stage area. Big monitors on the walls were showing the same "live" feed as they had five minutes ago, but now a red banner flashed at the bottom of the screens: ACTIVE SHOOTER AT COMPLETE PICTURES BUILDING.

Michaela pointed at a moving figure on the screen. "That looks like you—but—"

Marius nodded. "Uh-huh. Apparently I like assault rifles."

Adan, their cameraman, had called up a local news feed after the first shouts of panic and confusion filtered through the studio's thick doors. What it showed was entirely and completely not what the three of them were seeing. Marius was inside the windowless second-floor studio, empty-handed, yet the monitors showed what looked like a drone feed of him moving into and out of view through the building's windows on the 10th floor. He was armed, and every now and then he would pause and shoot, calmly and methodically.

Marius shook his head in disgust. "Hey, Adan, could you give me a hand with this?"

The cameraman was hunched over his laptop. "Sorry, gotta figure out who's hijacked our signal."

"The same people who own the SWAT team," said Marius. "But forget what I said. I think you'd better get out of here."

"Why?"

"Look." Marius pointed at the monitors. They were showing a jumble of witness cell-phone videos. "There!" A jiggly shot showed a man lying in a corridor, dead eyes staring upward, a dark stain on his chest.

"But ..." Adan gaped. "That's me."

"Yes. This scene's not real yet. Listen, Adan, I mean it: you need to leave. The SWAT team's not on their way here to save any of us. They're here to make sure that what's up there"— he pointed—"matches what's down here."

“The SWAT team’s not on their way here to save any of us. They’re here to make sure

Michaela stood up, staring at him. “So it’s real. The antilibrary is real.”

“And soon, the antinet. Michaela, I’m so sorry. I shouldn’t have said anything. I knew they might be watching. Figured they’d be mad if I revealed it, but I never imagined they’d do this.” With the door at least somewhat secure, he went to join Adan at the mixing console. “Any luck?”

Adan shook his head. “I don’t know whether they’re intercepting our feed in here, at the router, or somewhere outside.”

“All this footage,” said Michaela. “It’s being computer-generated in real time? Like you said—by an antilibrary?”

“Yeah.” He smiled ruefully at her. “You got more than you bargained for, I guess. Honestly, I was only going to talk about Augmented Manners. I guess you pushed my buttons, I—”

“That’s okay.” She glanced at the monitors, a little rueful herself. “Digging for the truth is what I do, or used to. Apparently I’m good at it ... What do we do now?”

The TVs showed a black armored personnel carrier plowing up the avenue, with the Complete Pictures building a few blocks ahead of it.

“I think,” Marius said with a grimace, “we’re about to disappear.”

For Michaela, this morning had promised something other than bitter disappointment, for a change. She’d spent the past two years trying to rebuild her career, after her exposé of campaign irregularities in the last election. She’d been sued for libel, doxed, and made a pariah by the winning side, and ultimately lost her job.

Now, after months of taping flower shows and costume conventions for local TV, she’d finally chased down a friend of a friend who knew someone and managed to net an exclusive interview with a rising IT god. Marius Rivas, already a minor celebrity in the hacker world, was young, charismatic, and engaging, not yet rich but certain to be a billionaire in a few years; and he was notoriously hard to find, much less interview. Michaela used his product, Augmented Manners. Hell, almost everybody she knew did—one in 10 American adults, according to the company’s publicity handout. On the street, people were already calling it “the new iPhone,” the indispensable invention of the decade.

But traditional media didn't know what to make of it. Their hesitation had given her the opening she needed.

"Welcome back to the Complete Pictures News Feed. I'm Michaela Kline—yes, *that* Michaela Kline—and today I'm talking to Marius Rivas, the founder and CEO of Augmented Manners." She turned to him with a smile. "Marius, you've called it 'the first true political app.' It's just a phone app, but it has this magic ability to get people working together. How does it do this?"

Marius grinned and shrugged, a *gee-whiz* gesture that was going to do great things for his image. "It's pretty straightforward, actually. You see, when people talk, they often mean entirely different things by the same words. Say, what I mean by 'liberal' may be very different from what you mean. Or, say, 'family values.' Or even 'fairness,' or 'truth.' And that's the problem: people who think they share the same language are talking past one another, because in fact they don't share an understanding of the part of the language that matters—the *connotations* of words.

"Augmented Manners uses simultaneous translation technology to translate *within* the English language, rather than between it and another one. If the app has access to the social-media profile of the person you're talking to, it can assess pretty well how they express themselves. Using your own voice, it will change trigger words and phrases into more neutral synonyms. Basically, it acts like a translator, or mediator, that helps the other side hear what you *mean* rather than just what you say."

that what's up there matches what's down here."

"Why do you think it's caught on so quickly?"

He started talking happily about adoption rates and structured dialogues, and Michaela began to relax.

And maybe it would have all turned out all right if she hadn't asked Marius where he'd come from.

"The Navy had its own cyber-warfare unit. I did everything I could to get myself transferred to it. I wanted to serve my country in the best way I knew how—by programming. The first project I got put on was for fighting fake news. We were told to figure out a way to defuse it—because it's propaganda, right? It's misinformation, or lies that serve the enemy. And I thought, Great! This is what I was meant to do.

"The cyber-war project focused on how to suppress the fake news and amplify the signal of the truth. There were fake-news factories, so one of my first jobs was to find ways to hack into them. I'm pretty good; it worked. We were able to shut down the big ones. But it made no difference. Fake news had gone viral. There were as many sources of it as there were people who doubted Big Media. A whole segment of the population was spamming the net—how do you fight that?

"I started thinking about it differently, focusing not on what fake news is but what it's for. What's its product? The answer is simple: doubt. Fake news is designed to sow doubt. I went to my team lead, a senior officer named Cather, and told him that playing whack-a-mole with the sources wasn't going to work. If fake news is a technology of doubt, we needed to build a technology of trust.

"Cather had a different plan. Why not beat them at their own game? The factories had minuscule resources compared with ours. We had supercomputers that could manufacture disinformation on the fly. Why not use them to spam the spammers?

"The project team was big. Maybe I should have noticed Cather's hiring practices sooner. I woke up one day and realized I was surrounded by a—well, a certain type of person ..."

"What type?" she asked.

"Intense, quiet young white men. Humorless. Even as civilians they'd buzz-cut their hair. Dress in black. Their e-mail signatures are quotes from Jordan Peterson."

"Oh," she said.

"The project was all about defanging fake news. The problem was, these guys couldn't actually see what was wrong with it to begin with. To them, there was no such thing as *real* news—never had been. There was only more or less powerful messaging. I'd gone along with Cather's plan to fight fire with fire because I still thought we were doing it to clear a path for truth. But inch by inch, almost unnoticeably, Cather and his boys steered the project away from using fake news to fight fake news ... toward just making better versions of it."

He couldn't seem to stop himself now; he talked in more and more detail about what he'd done, and as he did he got more and more passionate. Michaela was intensely aware of the cameras, and that everything he said was streaming live. He didn't seem to care—he'd fallen over some cliff of decision, and there was no going back for him.

"Cather rebranded us as Project Antilibrary. And an antilibrary is exactly what we built. We took the newest game engines, which can produce photorealistic video in real time, and we mated them to photo databases of places and people. At first, we used streets and boardrooms in Russia and China. It turned out to be so easy to build the databases, though, that we eventually said, 'What the hell—let's do the whole world!' We built a system that could deepfake live video from anywhere on the planet, of anybody we had a photo or video record of. We could deepfake voices, too—perfectly, of course.

"There were already systems to write fake scientific papers; we improved those so we could generate blog posts, TV news spots read out by America's favorite announcers, supporting documentation, you name it. I kind of surfed my way through

all this, morally. I rode a wave of denial, the way I guess concentration camp guards did.

"Those quiet young men, though ... They saw something in what we were building that I couldn't even guess at, at the time. By the time I left, we had a tool that could produce a literal library's worth of entirely bogus material: videos, articles, subtly nonsensical books cross-referenced and supported by other nonsensical books—a true antilibrary—and do it faster than the news cycle, or even social media, could keep up."

"So you left the Navy?"

He nodded. "I could see where Cather's work was leading. And I felt, morally, I had to counter it somehow. But we'd worked together, analyzing how disinformation operates. I knew there was no way to attack it directly, so I designed Augmented Manners.

"Can you tell me more about that? What was Augmented Manners supposed to do that would counter fake news? It seems like it's solving a different problem, the problem of how people communicate—"

"But that's exactly it. Fake news is designed to break down our ability to trust one another. So I built something that would help you see the real person you were talking to. It would be like your own personal diplomat. You go around, your diplomat talks to other people's diplomats. It finds common ground, builds bridges, learns what your values are and how to mesh what you want with what other people want. Yes, it really is the first true political app, in the sense that it does what politics is supposed to do: align everybody with collective decisions. It didn't try to attack the vector of the fake-news disease; it's designed to immunize us against its effects."

"And it works! Augmented Manners is the hottest app on three different platforms!"

Marius sighed and shook his head. To Michaela's alarm, he looked defeated rather than confident. "I've been having nightmares ever since we launched Augmented Manners ... I think they're my subconscious trying to get me to think it through, to recognize where Cather's really going with his work."

"Which is?" She glanced uneasily at Adan, who was frowning back from behind the mixing console.

Marius was sitting forward now, hands clasped between his knees. "The antilibrary runs on a cluster of supercomputers. But you know, computers are getting faster and cheaper all the time. And with neural-net and quantum machines, we knew we could double our speed, triple it. I understand the endgame now. It's to build a system that can tailor every piece of media a person consumes to whatever message you want to convey, on the fly and at any level of resolution. Not just a library's worth of bogus truth, but any amount, up to and including an internet's worth. I've started calling it the antinet.

"And here's the thing: he's built it."

That's when the monitors flickered and then showed Marius Rivas kicking back his chair and screaming as he brandished an assault rifle conjured from nowhere.

Marius watched the monitors. They showed police and SWAT pulling up to the building. He eyed the door he'd just barricaded. "Maybe I'm wrong. I think we need to leave."

"If you're right, what good will that do?" Adan slumped in his chair. "Even if we go online and wave our shorts at the world, they'll claim we're the fake news."

Marius stepped back. "That's true—right now. But you know what I said about not being able fight the antinet directly? I ... might be wrong about that."

Michaela blinked, then looked down from the monitors to him. "What do you mean?"

"When Cather and I were building the antilibrary, we were fixated on politics and society. We looked at how deepfaking could impact elections, opinions, social movements. If we're lucky, that's how Cather and his clients have continued to think."

"Why?"

"You know what groupthink is? When a whole team shares a worldview so intensely that they literally can't see things—obvious things to you or me—outside it? If we're lucky ... Put it this way: I was already mostly through the Manners project before some friends who aren't interested in politics and society pointed out a ... maybe an Achilles' heel to both the antilibrary and the antinet. Their solution is already out there. I just don't know whether it'll reveal itself in time to save us."

She blinked at him. "What are you talking about?"

He shrugged. "See, if you're purely political, you might think that there'd be no consequence to faking everybody's news and social media. You could just steer them like cattle, and that would be that.

"But the whole modern economy runs on accurate data. Without it, production targets are wrong, planes don't leave on time, things get shipped to wrong addresses. The margins of error are so fine now that, say, failing to predict the exact weather for New York in four days will have an international cascade effect."

He dug in his pocket for his phone, looked at it, and grimaced. "Do you, do either of you, have a signal?"

They both brought out their smartphones and held them up. They were deep inside the building. "No bars," said Adan.

"I've got one," said Michaela.

"Can I borrow that?" Marius reached for her phone. Opening its web browser, he tapped in an address. "It's a contradiction, right?" he muttered as he worked. "The system that's emerging, at one and the same time it wants to lie to us about everything, even while it's demanding perfectly accurate data at all scales for it to run at all. Everything from taxi locations to strain-gauge readings in bridge abutments—it all has to be accurate."

"They'll be here any second," said Michaela. "Are we getting out of this place or not?"

“My friends will be able to find you—and to hide you, It’s time to hand it over to someone like you who knows how to pursue the truth.”

He barely heard her. Moving around to look at the phone screen, she saw that he was filling out a long form on a gray web page that had no graphics or title. “Some coders came to me a while back,” he said. “They explained about the Internet of Things blockchain that all these sensors report to. Every sensor puts an encrypted signature on each packet of data it sends, and these are knitted together on a blockchain at the next level, which reports up the line to the next. It’s a chain of provenance, with auditable, publicly available proof at every step that the data hasn’t been tampered with.”

He laughed. “Imagine me when I realized it: we don’t need to build the anti-antinet! It already exists!”

“I don’t understand,” said Michaela hopelessly. “How does that help us?”

“Because if you can build an incorruptible pyramid of facts out of IoT sensor data,” said Marius, “you can do the same thing with journalistic facts.”

“Oh.” Her eyes went wide. “Oh?”

The door rattled as someone or something tried to force it open. “This is the police! Mr. Rivas, open the door and slide out your weapon. You have one minute to comply.”

“They’re not negotiating?” Michaela heard herself say. “Why aren’t they negotiating?”

“Quick!” Adan came out from behind the console and grabbed Michaela’s arm. He hissed under his breath: “I’ve pulled the shelves out of the storage cupboard under the mixer. Maybe you can cram yourself in there—”

“Good idea,” Marius whispered. Then he put his back against the wall, next to the door. “Hey!” he shouted. “How many people do you think are in here?”

There was silence from outside. “You have video of Adan Stokley being killed, right?” he continued. “Well—Adan, can you come over here?” Adan reluctantly went to join him. “Talk to them,” Marius hissed.

“Uh, hi. This is Adan. I’m not dead.”

“Who is this?” barked the voice from outside. “What are you playing at?”

“I’m not a hostage,” Adan shouted. “We’re fine!”

“These friends of mine, they told me they’ve been building a hybrid,” said Marius, speaking barely above a whisper so only

Michaela and Adan could hear. He finished entering his data and clicked the old-style OK button at the bottom of the form. Then he blew out a heavy sigh. "Their system lets anybody ask the IoT blockchain for its testimony as an incorruptible witness to anything, provided the data doesn't violate somebody's privacy. So, for instance, there are hundreds of diagnostic systems in this room—in the lights, the walls, even the chairs—that report wear and tear up the line to the manufacturers, building owners ... But that data can be used to testify about other things that're happening in the room. Or aren't happening. Like, from pressure-wave readings, whether someone has fired a gun in here."

"Get under here!" As Adan helped Michaela squeeze under the mixing console, she protested to Marius, "These friends of yours, they can prove what really happened here!"

Adan closed the cupboard. Marius came over to it and leaned in toward its door so Michaela could hear. "Once it's running, the testimonial service can't be shut down without crashing the economy. But it's not ready yet; we hadn't figured out how to make it accessible to everyone. Listen, just now I sent my friends a message. Told them to tune the system to track one particular set of inputs: yours, Michaela. My friends will be able to find you—and to hide you, I hope. I've given them as much help as I can. It's time to hand it over to someone like you who knows how to pursue the truth and tell the world about it."

The studio door crashed inward and the lights went out, and then it was all chaos and gunfire.

I hope. I've given them as much help as I can.

Michaela could barely breathe in the hot and claustrophobically dark cupboard. She heard somebody shout, "Anybody else?"

"Nothing on infrared," said another voice. Bodies rushed back and forth, and then suddenly everything was silent.

She was too terrified to move, sure she was about to be found or shot through the cupboard doors. The silence stretched out for long seconds, then a minute.

Light reappeared in the crack between the doors.

Michaela pushed, and tumbled out. She stood to find the studio empty, its door hanging open from one hinge.

The chairs were overturned; she moved past them to the door and looked cautiously through it. The outer office, which was usually bustling, was empty, its lights also off. They came on as she stepped out.

She frowned at the humming fluorescents and then went back for her phone, which was lying by one of the chairs. She grabbed her purse from the mixing console. After standing there for a long moment, she called up a local news feed on the phone. It wasn't hard to find the live feed she was after: the Complete Pictures building was front and center on the main page.

The SWAT team pushed their way out of the building as police built a cordon around the front entrance. In their wake, ambulance attendants were bringing out stretchers. There were three of them.

The sightless eyes of Marius Rivas panned past the camera's jittery perspective. His limp arm flopped over the side of the stretcher. His shirt was dark with blood. Behind him came Adan, his eyes closed, arms folded across his chest. Michaela sobbed at the sight.

And then her own eyes stared back at her from the third stretcher. For just a moment, she saw herself as others must see her, an image on a screen, the slightly crooked nose bent in the opposite direction from what she saw in the mirror, her hair and shoulders subtly different. Yet this Michaela Kline was clearly dead.

The stretchers moved on and the camera panned to a breathless commentator. Step by step, eyes not really tracking anything, Michaela drifted through the office to the stairs, and down them. Silence on the floor below. She continued past the main floor, headed for the parking garage.

Her car was one flight further down, but some instinct warned her not to go to it. SWAT clearly weren't watching the news. They hadn't yet spotted the mismatch in the number of stretchers coming out of the building, but that couldn't last long. On the landing where she stood was an emergency exit that led to an alley.

She was dead. Everybody knew it. She thought about her life of frustrations, about the empty apartment that only her landlady would care to clear out, about her dead-end job. Even if she cut some sort of deal with this Cather or his masters and miraculously came back to life, what was the point of trying to rebuild her career? It would be a lie, a life spent feeding the antilibrary.

She started to check the broadcast on her phone again, then remembered what Marius had been doing on it. She dug in her purse for a pen and notebook, and wrote down the numbered web address of the site he'd been on. Then she battered the phone until it broke, and pushed open the door to the alley.

There was nobody outside, just drifting plastic bags and the blue hulk of a dumpster. The murmur of a crowd came from the right; Michaela turned left and started walking.

All she had now was an address, but she had been an investigative reporter. There really were facts, there really was a chain of truths, and she would follow that chain—

Until it led her out of the antilibrary.

Karl Schroeder is a futurist and the author of 11 science fiction books. His latest novel is The Million. His next novel, due in 2019, is Stealing Worlds, about the future of work (and burglary) in a near-future, "post-real" America.

How to tell if you're arguing with a bot

By Sarah Cooper

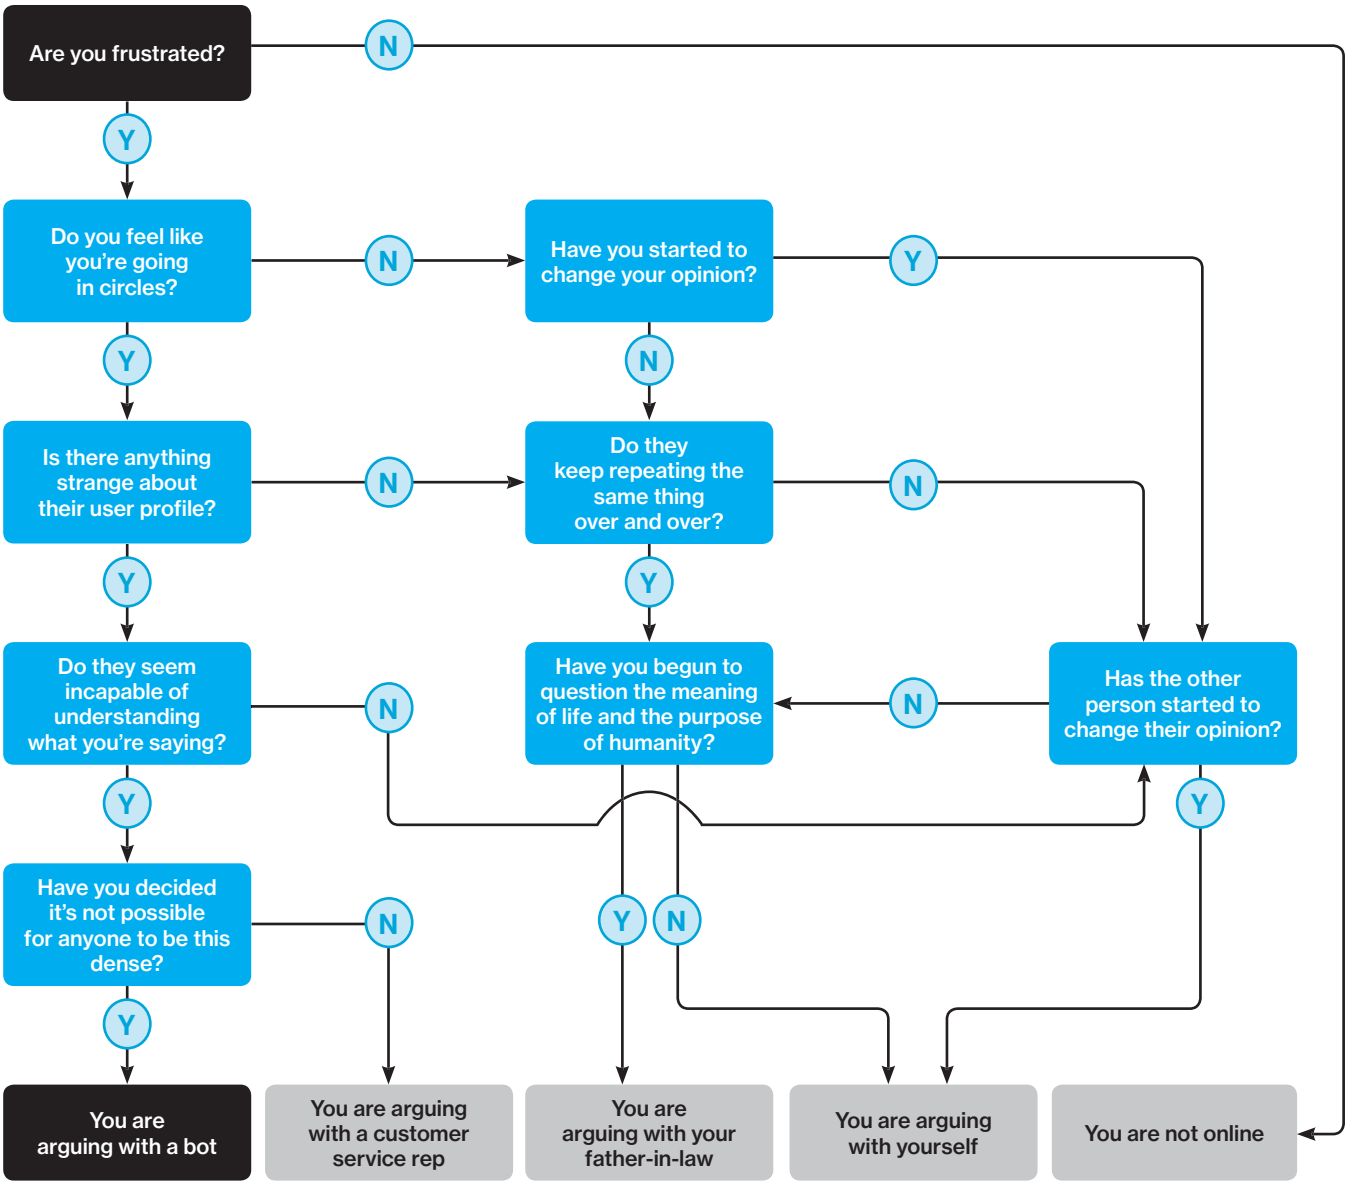

Sarah Cooper is a writer, comedian, and creator of the satirical blog TheCooperReview.com.

MIT Technology Review (ISSN 1099-274X), September/October 2018 issue, Reg. US Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Entire contents ©2018. The editors seek diverse views, and authors' opinions do not represent the official policies of their institutions or those of MIT. Periodicals postage paid at Boston, MA, and additional mailing offices. Postmaster: Send address changes to MIT Technology Review, Subscriber Services, PO Box 5001, Big Sandy, TX 75755, or via the internet at www.technologyreview.com/customerservice. Basic subscription rates: $59.94 per year within the United States; in all other countries, US$69. Publication Mail Agreement Number 40621028. Send undeliverable Canadian copies to PO Box 1051, Fort Erie, ON L2A 6C7. Printed in USA. Audited by the Alliance for Audited Media.


Blockchain.
Hype? Hope?
The future
is in between.
The future
is here.
BUSINESS
OF BLOCKCHAIN
MIT Technology Review
May 2, 2019 MIT Media Lab Cambridge, MA
technologyreview.com/blockchain2019


Creating new technologies
Informing public policy
Eradicating deadly disease
Improving patient outcomes
Keeping airports safe
BATTELLE
battelle.org/tech
It can be done